Package 'ergmito'

December 20, 2025

Version 0.3-2

Title Exponential Random Graph Models for Small Networks

Description Simulation and estimation of Exponential Random Graph Models (ERGMs)
for small networks using exact statistics as shown in Vega Yon et al. (2020)
<DOI:10.1016/j.socnet.2020.07.005>. As a difference from the 'ergm'
package, 'ergmito' circumvents using Markov-Chain Maximum Likelihood Estimator
(MC-MLE) and instead uses Maximum Likelihood Estimator (MLE) to fit ERGMs
for small networks. As exhaustive enumeration is computationally feasible for
small networks, this R package takes advantage of this and provides tools for
calculating likelihood functions, and other relevant functions, directly,
meaning that in many cases both estimation and simulation of ERGMs for
small networks can be faster and more accurate than simulation-based
algorithms.

Depends R (>= 3.3.0)

RoxygenNote 7.3.3

Encoding UTF-8

Imports ergm, network, MASS, Rcpp, texreg, stats, parallel, utils,
methods, graphics

LinkingTo Rcpp, RcppArmadillo

License MIT + file LICENSE

Suggests covr, sna, lmtest, fmcmc, coda, knitr, rmarkdown, tinytest

VignetteBuilder knitr

LazyData true

URL https://muriteams.github.io/ergmito/

BugReports https://github.com/muriteams/ergmito/issues

Language en-US

NeedsCompilation yes

Author George Vega Yon [cre, aut] (ORCID:
<https://orcid.org/0000-0002-3171-0844>),
Kayla de la Haye [ths] (ORCID: <https://orcid.org/0000-0002-2536-7701>),
Army Research Laboratory and the U.S. Army Research Office [fnd] (Grant
Number W911NF-15-1-0577)

Maintainer George Vega Yon <g.vegayon@gmail.com>

Repository CRAN

Date/Publication 2025-12-19 23:30:02 UTC

Contents

---

as_adjmat   *An alternative to* as.matrix *to retrieve adjacency matrix fast*

Description

This function does not perform significant checks. Furthermore, this function won't keep the
row/col names.

Usage

```
as_adjmat(x)
```

Arguments

x   An object to be coerced as an adjacency matrix.

Examples

```
set.seed(1231)
x <- matrix_to_network(rbernoulli(rep(5, 100)))
benchmarkito(
  as_adjmat = as_adjmat(x),
  as.matrix = lapply(x, as.matrix)
)
```

---

benchmarkito   *Utility to benchmark expression in R*

Description

This is just an internal utility included in the package which is not designed to be accurate. If you need accurate benchmarks, you should take a look at the **microbenchmark** and **bench** R packages.

Usage

```
benchmarkito(..., times = 100, rand_ord = TRUE)
```

Arguments

...	List of expressions to benchmark.
times	Integer scalar. Number of replicates.
rand_ord	Logical. When TRUE, the expressions are executed in a random order.

Details

The print method prints a summary including quantiles, relative elapsed times, and the order (fastest to slowest).

Value

A data frame of class ergmito_benchmark with times rows and the following columns:

- id Integer. Id of the expression.
- expr Factor. Expression executed.
- user.self, sys.self, elapsed, sys.child time in seconds (see proc.time()).

Includes the following attributes: ncalls, call, label, and expr.

Examples

```
bm <- benchmarkito(
  exp(1:100000),
  sqrt(1:100000),
  times = 20
)

plot(bm)
print(bm)
```

---

blockdiagonalize *Block-diagonal models using* ergm

Description

These two functions are used to go back and forth from a pooled ergm vs a blockdiagonal model, the latter to be fitted using ergm::ergm.

Usage

```
blockdiagonalize(x, attrname = "block")

splitnetwork(x, attrname)

ergm_blockdiag(formula, ...)
```

Arguments

x	In the case of blockdiagonalize, a list of networks or matrices. For splitnetwork a single network object with a vertex attribute that can be used to split the data.
attrname	Name of the attribute that holds the block ids.
formula	An ergm model which networks' will be wrapped with blockdiagonalize (see details).
...	Further arguments passed to the method.

Details

The function ergm_blockdiag is a wrapper function that takes the model's network, stacks the networks into a single block diagonal net, and calls ergm::ergm with the option constraints = blockdiag("block").

One side effect of this function is that it loads the ergm package via requireNamespace, so after executing the function ergm the package will be loaded.

Value

An object of class ergm::ergm.

Examples

```
library(ergm)
data(fivenets)

fivenets2 <- blockdiagonalize(fivenets, attrname = "block") # A network with
ans0 <- ergm(
  fivenets2 ~ edges + nodematch("female"),
  constraints = ~blockdiag("block")
  )
ans1 <- ergmito(fivenets ~ edges + nodematch("female"))

# This is equivalent
ans2 <- ergm_blockdiag(fivenets ~ edges + nodematch("female"))
```

---

check_support   *Check the convergence of ergmito estimates*

Description

This is an internal function used to check the convergence of the optim function.

Usage

```
check_support(target_stats, stats_statmat, threshold = 0.8, warn = TRUE)

check_convergence(optim_output, model, support, crit = 5)
```

Arguments

target_stats, stats_statmat

> See [ergmito_formulae](#).

threshold   Numeric scalar. Confidence range for flagging an observed statistic as potentially near the boundary.

warn   logical scalar.

optim_output   A list output from the [stats::optim](#) function.

model   An object of class [ergmito_loglik](#).

support   As returned by check_support.

crit   Numeric scalar. Level at which a parameter estimate will be questioned.

Value

A list with the following components:

- par Updated set of parameters
- vcov Updated variance-covariance matrix

- valid Vector of integers with the parameters that are marked as OK.
- status Return code of the analysis. See details.
- note A note describing the status.

Return codes

The function makes an analysis of the outcome of the model and makes the corresponding adjustments when required. In particular, we check:

1. Whether the optimization algorithm converged or not

2. If the obtained estimates maximize the function. If this is not the case, the function checks whether the MLE may not exist. This usually happens when the log-likelihood function can improve by making increments to parameters that are already tagged as large. If the ll improves, then the value is replaced with Inf (+- depending on the sign of the parameter).

3. If the Hessian is semi-positive-definite, i.e. if it is invertible. If it is not, it usually means that the function did not converged, in which case we will use MASS::ginv instead.

The return codes are composed of two numbers, the first number gives information regarding of the parameter estimates, while the second number give information about the variance-covariance matrix.

Column 1:

- 0: Converged and estimates at the max.
- 1: It did not converged, but I see no issue in the max.
- 2: One or more estimates went to +/-Inf
- 3: All went to hell. All estimates went to +/-Inf

Column 2:

- 0: Hessian is p.s.d.
- 1: Hessian is not not p.s.d.

Possible codes and corresponding messages:

- 00 All OK (no message).
- 01 optim converged, but the Hessian is not p.s.d..
- 10 optim did not converged, but the estimates look OK..
- 11 optim did not converged, and the Hessian is not p.s.d..
- 20 A subset of the parameters estimates was replaced with +/-Inf..
- 21 A subset of the parameters estimates was replaced with +/-Inf, and the Hessian matrix is not p.s.d..
- 30 All parameters went to +/-Inf suggesting that the MLE may not exists..

---

count_stats   *Count Network Statistics*

Description

This function is similar to what ergm::summary_formula does, but it provides a fast wrapper suited for matrix class objects (see benchmark in the examples).

Usage

```
count_stats(X, ...)

AVAILABLE_STATS()

## S3 method for class 'formula'
count_stats(X, ...)

## S3 method for class 'list'
count_stats(X, terms, attrs = NULL, ...)
```

Arguments

X	List of square matrices. (networks)
...	Passed to the method.
terms	Character vector with the names of the statistics to calculate. Currently, the only available statistics are: 'mutual', 'edges', 'ttriad', 'ctriad', 'ctriple', 'node-icov', 'nodeocov', 'nodematch', 'triangle', 'balance', 't300', 't102', 'absdiff', 'idegree1.5', 'odegree1.5', 'ostar1', 'ostar2', 'ostar3', 'ostar4', 'istar1', 'istar2', 'istar3', 'istar4'.
attrs	A list of vectors. This is used when term has a nodal attribute such as nodeicov(attrname="").

Value

A matrix of size length(X) * length(terms) with the corresponding counts of statistics.

Examples

```
# DGP
set.seed(123199)
x <- rbernoulli(rep(5, 10))
ans0 <- count_stats(x, c("mutual", "edges"))

# Calculating using summary_formula
ans1 <- lapply(x, function(i) {
  ergm::summary_formula(i ~ mutual + edges)
})
```

```
ans1 <- do.call(rbind, ans1)

# Comparing
all.equal(unname(ans0), unname(ans1))

# count_stats is vectorized (and so faster)
bm <- benchmarkito(
  count_stats = count_stats(x, c("mutual", "edges")),
  lapply      = lapply(x, function(i) {
  ergm::summary_formula(i ~ mutual + edges)
}), times = 50
)

plot(bm)
```

---

ergmito_boot   *Bootstrap of ergmito*

Description

Bootstrap of ergmito

Usage

```
ergmito_boot(x, ..., R, ncpus = 1L, cl = NULL)
```

Arguments

x	Either a formula or an object of class ergmito.
...	Additional arguments passed to the method.
R	Integer. Number of replicates
ncpus	Integer Number of CPUs to use. Only recommended if ergmito was not compiled with OpenMP (otherwise it will be slower).
cl	An object of class cluster (see makePSOCKcluster)

Details

The resulting sample of parameters estimates is then used to compute the variance-covariance matrix of the model. Cases in which Inf/NaN/NA values were returned are excluded from the calculation.

Value

An object of class ergmito_boot and [ergmito](). This adds three elements to the ergmito object:

- R The number of replicates.
- sample A vector of length R with the cases used in each replicate.
- dist The distribution of fitted parameters.
- nvalid the number of cases used for computing the covar.
- timer_boot records the time the whole process took.

Examples

```
data(fivenets)
set.seed(123)
ans0 <- ergmito(fivenets ~ edges + ttriad)
ans1 <- suppressWarnings(ergmito_boot(ans0, R = 100))
ans2 <- suppressWarnings(ergmito_boot(fivenets ~ edges + ttriad, R = 100))

# Checking the differences
summary(ans0)
summary(ans1)
summary(ans2)
```

---

ergmito_formulae   *Processing formulas in* ergmito

Description

Analyze formula objects returning the matrices of weights and sufficient statistics to be used in the model together with the log-likelihood and gradient functions for joint models.

Usage

```
ergmito_formulae(
  model,
  model_update = NULL,
  target_stats = NULL,
  stats_weights = NULL,
  stats_statmat = NULL,
  target_offset = NULL,
  stats_offset = NULL,
  env = parent.frame(),
  ...
)
```

Arguments

model	A formula. The left-hand-side can be either a small network, or a list of networks.
model_update	A formula. If specified, the after computing the sufficient statistics (observed and support), the model is updated using stats::model.frame(). This includes processing offset terms.
target_stats	Observed statistics. If multiple networks, then a list, otherwise a named vector (see ergm::summary_formula).
stats_weights, stats_statmat	
	Lists of sufficient statistics and their respective weights.
target_offset, stats_offset	
	See exact_loglik().
env	Environment in which model should be evaluated.
...	Further arguments passed to ergm::ergm.allstats.

Details

One of the main advantages of been able to compute exact likelihoods is that we can build arbitrarily complex models in the same way that we would do in the context of Generalized Linear Models, this is, adding offset terms, interaction effects, or transformations of statistics without much effort.

In particular, if the user passes a formula via model_update, the cannonical additive ERGM can be modified to include other terms, for example, if we wanted to add an interaction effect of the nodematch("age") with network size, we can simply type

```
model_update = ~ . + I(nodematch.age * n)
```

The I() function allows operating over variables in the model, in this case, we took the nodematch.age variable (which is the name that ergm::ergm() assigns to it after computing the sufficient statistics) and multiplied it by n, which is the network size (this variable is included by default).

By default, the ergm package calculates up to 2^16 unique values for the vector of sufficient statistics. This results in issues if the user tries to fit a model with too heterogenous networks or sets of attributes. To deal with this it suffices with adding the option maxNumChangeStatVectors in the ergmito call, e.g.:

```
# Networks of size 5 have up to 2^20 unique sets of sufficient statistics
ergmito(..., maxNumChangeStatVectors = 2^20)
```

See more in ?ergm::ergm.allstats.

Value

A list of class ergmito_loglik.

- loglik A function. The log-likelihood function.
- grad A function. The gradient of the model.

- stats_weights,stats_statmat two list of objects as returned by ergm::ergm.allstats.

- target_offset,stats_offset A vector of offset terms and a list of vectors of offset terms, one for the target stats and the other for the support of the suffient statistics (defaults to 0).

- model A formula. The model passed.

- npars Integer. Number of parameters.

- nnets Integer. Number of networks in the model.

- vertex_attr Character vector. Vertex attributes used in the model.

- term_names Names of the terms used in the model.

Examples

```
data(fivenets)
model0 <- ergmito_formulae(fivenets ~ edges + nodematch("female"))
print(model0)
model0$loglik(c(-2, 2))

# Model with interaction effects and an offset term
model1 <- ergmito_formulae(
  fivenets ~ edges + nodematch("female"),
  model_update = ~ . + offset(edges) + I(edges * nodematch.female)
)
```

---

ergmito_gof   *Goodness of Fit diagnostics for ERGMito models*

Description

Goodness of Fit diagnostics for ERGMito models

Usage

```
gof_ergmito(
  object,
  GOF = NULL,
  GOF_update = NULL,
  probs = c(0.05, 0.95),
  sim_ci = FALSE,
  R = 50000L,
  ncores = 1L,
  ...
)

## S3 method for class 'ergmito_gof'
plot(
  x,
  y = NULL,
```

```
    main = NULL,
    sub = NULL,
    tnames = NULL,
    sort_by_ci = FALSE,
    ...
)
```

Arguments

`object`	An object of class ergmito.
`GOF`	Formula. Additional set of parameters to perform the GOF.
`GOF_update`	Formula. See the section on model updating in `ergmito_formulae()`.
`probs`	Numeric vector. Quantiles to plot (see details).
`sim_ci`	Logical scalar. If `FALSE`, the default, it will compute the quantiles analytically, otherwise it samples from the ERGM distribution.
`R`	Integer scalar. Number of simulations to generate (passed to sample). This is only used if `sim_ci = TRUE`.
`ncores`	Integer scalar. Number of cores to use for parallel computations (currently ignored).
`...`	Further arguments passed to stats::quantile.
`x`	An object of class ergmito_gof.
`y`	Ignored.
`main, sub`	Title and subtitle of the plot (see graphics::title).
`tnames`	A named character vector. Alternative names for the terms.
`sort_by_ci`	Logical scalar. When `TRUE` it will sort the x-axis by the with of the CI in for the first parameter of the model.

Details

The Goodness of Fit function uses the fitted ERGMito to calculate a given confidence interval for a set of sufficient statistics. By default (and currently the only available option), this is done on the sufficient statistics specified in the model.

In detail, the algorithm is executed as follow:

For every network in the list of networks do:

1. Calculate the probability of observing each possible graph in its support using the fitted model.

2. If `sim_ci = TRUE`, draw R samples from each set of parameters using the probabilities computed. Then using the `quantile` function, calculate the desired quantiles of the sufficient statistics. Otherwise, compute the quantiles using the analytic quantiles using the full distribution.'

The plot method is particularly convenient since it graphically shows whether the target statistics of the model (observed statistics) fall within the simulated range.

The print method tries to copy (explicitly) the print method of the `gof` function from the `ergm` R package.

Value

An object of class `ergmito_gof`. This is a list with the following components:

- `ci` A list of matrices of length `nnets(object)` with the corresponding confidence intervals for the statistics of the model.
- `target_stats` A matrix of the target statistics.
- `ergmito.probs` A list of numeric vectors of length `nnets(object)` with the probabilities associated to each possible structure of network.
- `probs` The value passed via `probs`.
- `model` The fitted model.
- `term_names` Character vector. Names of the terms used in the model.
- `quantile.args` A list of the values passed via `...`.

Examples

```
# Fitting the fivenets model
data(fivenets, package = "ergmito")
fit <- ergmito(fivenets ~ edges + nodematch("female"))

# Calculating the gof
ans <- gof_ergmito(fit)

# Looking at the results
ans
plot(ans)
```

---

exact_loglik   *Vectorized calculation of ERGM exact log-likelihood*

Description

This function can be compared to ergm::ergm.exact with the statistics not centered at `x`, the vector of observed statistics.

Usage

```
exact_loglik(x, params, ..., as_prob = FALSE)

## Default S3 method:
exact_loglik(
  x,
  params,
  stats_weights,
  stats_statmat,
  target_offset = double(nrow(x)),
  stats_offset = lapply(stats_weights, function(i) double(length(i))),
```

```
  ...,
  as_prob = FALSE
)

exact_gradient(x, params, ...)

## Default S3 method:
exact_gradient(
  x,
  params,
  stats_weights,
  stats_statmat,
  target_offset = double(nrow(x)),
  stats_offset = lapply(stats_weights, function(i) double(length(i))),
  ...
)

exact_hessian(
  params,
  stats_weights,
  stats_statmat,
  stats_offset = lapply(stats_weights, function(i) double(length(i)))
)
```

Arguments

x	Matrix. Observed statistics
params	Numeric vector. Parameter values of the model.
...	Arguments passed to the default methods.
as_prob	Logical scalar. When `TRUE`, the function returns probabilities instead of log-likelihoods.
stats_weights	Either an integer vector or a list of integer vectors (see [exact_loglik](#)).
stats_statmat	Either a matrix or a list of matrices (see [exact_loglik](#)).
target_offset	Numeric vector of length `nrow(target_stats)`.
stats_offset	List of numeric vectors, each of length equal to the lengths of vectors in `stats_weights` (see details).

Sufficient statistics

One of the most important components of `ergmito` is calculating the full support of the model's sufficient statistics. Right now, the package uses the function [ergm::ergm.allstats](#) which returns a list of two objects:

- `weights`: An integer vector of counts.
- `statmat`: A numeric matrix with the rows as unique vectors of sufficient statistics.

Since `ergmito` can vectorize operations, in order to specify weights and statistics matrices for each network in the model, the user needs to pass two lists `stats_weights` and `stats_statmat`. While both lists have to have the same length (since its elements are matched), this needs not to be the case with the networks, as the user can specify a single set of weights and statistics that will be recycled (smartly).

In the case of offset terms, these can be passed directly via `target_offset` and `stats_offset`. The first is a numeric vector of length equal to the number of networks in the model. The later is a list of vectors that is matched against `stats_weights`, so each of it's elements must be a numeric vector of the same length that in the list of weights. By default the offset terms are set to equal zero.

Examples

```
data(fivenets)
ans <- ergmito(fivenets ~ edges + nodematch("female"))

# This computes the likelihood for all the networks independently
with(ans$formulae, {
  exact_loglik(
    x       = target_stats,
    params = coef(ans),
    stats_weights = stats_weights,
    stats_statmat = stats_statmat
  )
})

# This should be close to zero
with(ans$formulae, {
  exact_gradient(
    x       = target_stats,
    params = coef(ans),
    stats_weights = stats_weights,
    stats_statmat = stats_statmat
  )
})

# Finally, the hessian
with(ans$formulae, {
  exact_hessian(
    params = coef(ans),
    stats_weights = stats_weights,
    stats_statmat = stats_statmat
  )
})
```

---

extract.ergmito   *Extract function to be used with the* texreg *package.*

Description

To be used with the **texreg** package. This function can be used to generate nice looking tables of ERGMitos estimates.

Usage

```
extract.ergmito(
  model,
  include.aic = TRUE,
  include.bic = TRUE,
  include.loglik = TRUE,
  include.nnets = TRUE,
  include.offset = TRUE,
  include.convergence = TRUE,
  include.timing = TRUE,
  ...
)
```

Arguments

model   An object of class `ergmito`.

`include.aic`, `include.bic`, `include.loglik`
                  See texreg::extract.

`include.nnets`   Logical. When true, it adds the Number of networks used to the list of gof statistics. This can be useful when running pooled models.

`include.offset` Logical. When equal to `TRUE`, it adds one line per offset term to the table, omitting sd and significance.

`include.convergence`
                  Logical. When true it, adds the convergence value of the stats::optim function (0 means convergence).

`include.timing` Logical, When true it will report the elapsed time in seconds.

`...`   Further arguments passed to the `base::summary()` of ergmito.

Examples

```
library(texreg)
data(fivenets)
ans <- ergmito(fivenets ~ edges + nodematch("female"))
screenreg(ans)
```

---

fivenets   *Example of a group of small networks*

Description

This list of networks was generated using the [new_rergmito](#) sampler from a set of 5 baseline networks with a random vector of female. The sufficient statistics that generate this data are edges and nodematch("female") with parameters -2.0 and 2.0 respectively.

Usage

```
fivenets
```

Format

An object of class list of length 5.

---

geodesic   *Geodesic distance matrix (all pairs)*

Description

Calculates the shortest path between all pairs of vertices in a network. This uses the power matrices to do so, which makes it efficient only for small networks.

Usage

```
geodesic(x, force = FALSE, ...)

geodesita(x, force = FALSE, ...)

## S3 method for class 'matrix'
geodesic(x, force = FALSE, simplify = FALSE, ...)

## S3 method for class 'network'
geodesic(x, force = FALSE, simplify = FALSE, ...)
```

Arguments

x	Either a list of networks (or square integer matrices), an integer matrix, a network, or an ergmito.
force	Logical scalar. If force = FALSE (the default) and nvertex(x) > 100 it returns with an error. To force computation use force = TRUE.
...	Further arguments passed to the method.
simplify	Logical scalar. When TRUE it returns a matrix, otherwise, a list of length nnets(x).

Examples

```
data(fivenets)
geodesic(fivenets)

# Comparing with sna
if (require("sna")) {
  net0 <- fivenets[[1]]
  net  <- network::network(fivenets[[1]])
  benchmarkito(
    ergmito = ergmito::geodesic(net0),
    sna     = sna::geodist(net), times = 1000
  )
}
```

---

induced_submat   *Extract a submatrix from a network*

Description

This is similar to [network::get.inducedSubgraph.](#) The main difference is that the resulting object
will always be a list of matrices, and it is vectorized.

Usage

```
induced_submat(x, v, ...)
```

Arguments

x	Either a list or single matrices or network objects.
v	Either a list or a single integer vector of vertices to subset.
...	Currently ignored.

Details

Depending on the lengths of x and v, the function can take the following strategies:

- If both are of the same size, then it will match the networks and the vector of indices.
- If length(x) == 1, then it will use that single network as a baseline for generating the sub-
  graphs.
- If length(v) == 1, then it will generate the subgraph using the same set of vertices for each
  network.
- If both have more than one element, but different sizes, then the function returns with an error.

Value

A list of matrices as a result of the subsetting.

Examples

```
x <- rbernoulli(100)
induced_submat(x, c(1, 10, 30:50))

x <- rbernoulli(c(20, 20))
induced_submat(x, c(1:10))
```

---

matrix_to_network   *Manipulation of network objects*

Description

This function implements a vectorized version of [network::network](.)`.adjmat`. It allows us to turn regular matrices into network objects quickly.

Usage

```
matrix_to_network(x, ...)

## S3 method for class 'matrix'
matrix_to_network(
  x,
  directed = rep(TRUE, length(x)),
  hyper = rep(FALSE, length(x)),
  loops = rep(FALSE, length(x)),
  multiple = rep(FALSE, length(x)),
  bipartite = rep(FALSE, length(x)),
  ...
)
```

Arguments

x	Either a single square matrix (adjacency matrix), or a list of these.
...	Further arguments passed to the method.
directed	Logical scalar, if `FALSE` then the function only checks the upper diagonal of the matrix assuming it is undirected.
hyper, multiple, bipartite	
	Currently Ignored. Right now all the network objects created by this function set these parameters as `FALSE`.
loops	Logical scalar. When `FALSE` (default) it will skip the diagonal of the adjacency matrix.

Details

This version does not support adding the name parameter yet. The function in the network package includes the name of the vertices as an attribute.

Just like in the network function, NA are checked and added accordingly, i.e. if there is an NA in the matrix, then the value is recorded as a missing edge.

Value

An object of class network. This is a list with the following elements:

- mel *Master Edge List*: A named list with length equal to the number of edges in the network. The list itself has 3 elements: inl (tail), outl (head), and atl (attribute). By default atl, a list itself, has a single element: na.
- gal *Graph Attributes List*: a named list with the following elements:
  - n Number of nodes
  - mnext Number of edges + 1
  - directed,hyper,loops,multiple,bipartite The arguments passed to the function.
- val *Vertex Attributes List*
- iel *In Edgest List*
- oel *Out Edgest List*

Examples

```
set.seed(155)
adjmats  <- rbernoulli(rep(5, 20))
networks <- matrix_to_network(adjmats)
```

---

new_ergmito_ptr   *Creates a new* ergmito_ptr

Description

After calculating the support of the sufficient statistics, the second most computationally expensive task is computing log-likelihoods, Gradients, and Hessian matrices of ERGMs. This function creates a pointer to an underlying class that is optimized to improve memory allocation and save computation time when possible.

Usage

```
new_ergmito_ptr(
  target_stats,
  stats_weights,
  stats_statmat,
  target_offset,
  stats_offset
)
```

Arguments

```
target_stats,     stats_weights,     stats_statmat,     target_offset,
stats_offset
```
                 see exact_loglik.

Details

This function is for internal used only. Non-advanced users are not encouraged to use it. See ergmito_formulae and exact_loglik for user friendly wrappers of this function.

Recycling computations

Some components of the likelihood, its gradient, and hessian can be pre-computed and recycled when needed. For example, it is usually the case that in optimization gradients are computed using a current state of the model's parameter, which implies that the normalizing constant and some other matrix products will be the same between the log-likelihood and the gradient. Because of this, the underlying class ergmito_ptr will only re-calculate these shared components if the parameter used changes as well. This saves a significant amount of computation time.

Scope of the class methods

To save space, the class creates pointers to the matrices of sufficient statistics that the model uses. This means that once these objects are deleted the log-likelihood and the gradient functions become invalid from the computational point of view.

---

new_rergmito   *ERGMito sampler*

Description

Create a sampler object that allows you simulating streams of small networks fast.

Usage

```
new_rergmito(model, theta, ...)

## S3 method for class 'ergmito_sampler'
x[i, ...]
```

Arguments

model	A formula.
theta	Named vector. Model parameters.
...	Further arguments passed to ergmito_formulae().
x	An object of class ergmito_sampler.
i	i is an integer vector indicating the indexes of the networks to draw.

Details

While the **ergm** package is very efficient, it was not built to do some of the computations required in the ergmito package. This translates in having some of the functions of the package (ergm) with poor speed performance. This led us to "reinvent the wheel" in some cases to speed things up, this includes calculating observed statistics in a list of networks.

The indexing method, `[.ergmito_sampler`, allows extracting networks directly by passing indexes. `i` indicates the index of the networks to draw, which go from 1 through `2^(n*(n-1))` if directed and `2^(n*(n-1)/2)` if undirected .

Value

An environment with the following objects:

- `calc_prob` A function to calculate each graph's probability under the specified model.
- `call` A language object with the call.
- `counts` A list with 3 elements: `stats` the sufficient statistics of each network, `weights` and `statmat` the overall matrices of sufficient statistics used to compute the likelihood.
- `network` The baseline network used to either fit the model or obtain attributes.
- `networks` A list with the actual sample space of networks.
- `probabilities` A numeric vector with each graph's probability.
- `sample` A function to draw samples. `n` specifies the number of samples to draw and `theta` the parameter to use to calculate the likelihoods.
- `theta` Named numeric vector with the current values of the model parameters.

The indexing method `[.ergmito_sampler` returns a list of networks

Examples

```
# We can generate a sampler from a random graph
set.seed(7131)
ans <- new_rergmito(rbernoulli(4) ~ edges, theta = -.5)

# Five samples
ans$sample(5)

# or we can use some nodal data:
data(fivenets)
ans <- new_rergmito(
  fivenets[[3]] ~ edges + nodematch("female"),
  theta = c(-1, 1)
)

# Five samples
ans$sample(5) # All these networks have a "female" vertex attr
```

---

nvertex   *Utility functions to query network dimensions*

Description

Utility functions to query network dimensions

Usage

```
nvertex(x)

nedges(x, ...)

nnets(x)

is_directed(x, check_type = FALSE)
```

Arguments

x	Either an object of class [ergmito](#), [network::network()](#), [stats::formula()](#), or [base::matrix()](#).
...	Further arguments passed to the method. Currently only nedges.network receives arguments (see [network::network.edgecount](#)).
check_type	Logical scalar. When checking for whether the network is directed or not, we can ask the function to return with an error if what we are checking is not an object of class network, otherwise it simply returns false.

Value

is_directed checks whether the passed networks are directed using the function [is.directed](#). In the case of multiple networks, the function returns a logical vector. Only objects of class network can be checked, otherwise, if check_type = FALSE, the function returns TRUE by default.

Examples

```
set.seed(771)
net <- lapply(rbernoulli(c(4, 4)), network::network, directed = FALSE)
is_directed(net)
is_directed(net[[1]])
is_directed(net ~ edges)
## Not run:
  is_directed(net[[1]][1:4, 1:4], check_type = TRUE) # Error

## End(Not run)
is_directed(net[[1]][1:4, 1:4])
```

---

plot.ergmito   *Function to visualize the optimization surface*

Description

General diagnostics function. This function allows to visualize the surface to be maximize at around a particular point.

Usage

```
## S3 method for class 'ergmito'
plot(
  x,
  y = NULL,
  domain = NULL,
  plot. = TRUE,
  par_args = list(),
  image_args = list(),
  breaks = 20L,
  extension = 4L,
  params_labs = stats::setNames(names(coef(x)), names(coef(x))),
  ...
)
```

Arguments

x	An object of class [ergmito](#).
y, ...	Ignored.
domain	A list.
plot.	Logical. When TRUE (default), the function will call [graphics::image](#) and plot all possible combination of parameters.
par_args	Further arguments to be passed to [graphics::par](#)
image_args	Further arguments to be passed to [graphics::image](#)
breaks	Integer scalar. Number of splits per dimension.
extension	Numeric. Range value of the function.
params_labs	Named vector. Alternative labels for the parameters. It should be in the form of c("orignial name" = "new name").

Details

It calculates the surface coordinates for each pair of parameters included in the ERGMito.

Value

A list of length choose(length(object$coef), 2) (all possible combinations of pairs of parameters), each with the following elements:

- z A matrix
- z A vector
- y A vector
- xlab A string. Name of the ERGM parameter in the x-axis.
- ylab A string. Name of the ERGM parameter in the y-axis.

The list is returned invisible.

See Also

The ergmito function.

Examples

```
set.seed(12)
x <- rbernoulli(c(4, 4, 5))

ans <- ergmito(x ~ edges + balance)

plot(ans)
```

---

powerset	*Power set of Graphs of size* n

Description

Generates the set of all possible binary networks of size n.

Usage

```
powerset(n, directed = TRUE, force = FALSE, chunk_size = 2e+05)
```

Arguments

n	Integer. Number of edges.
directed	Logical scalar. Whether to generate the power set of directed or undirected graphs,
force	Logical scalar. When TRUE it generates the power set for n>5, otherwise it returns with error.
chunk_size	Number of matrices to process at a time. If n = 5, then stack memory on the computer may overflow if chunk_size is relatively large.

Examples

```
powerset(2)
powerset(4, directed = FALSE)
```

---

predict.ergmito *Prediction method for* ergmito *objects*

Description

Takes an ergmito object and makes prediction at tie level. See details for information regarding its implementation.

Usage

```
## S3 method for class 'ergmito'
predict(object, newdata = NULL, ...)
```

Arguments

object	An object of class ergmito.
newdata	New set of networks (or network) on which to make the prediction.
...	Passed to new_rergmito, the workhorse of this method.

Details

After fitting a model with a small network (or a set of them), we can use the parameter estimates to calculate the likelihood of observing any given tie in the network, this is, the marginal probabilites at the tie level.

In particular, the function takes the full set of networks on the support of the model and adds them up weighted by the probability of observing them as predicted by the ERGM, formally:

$$\hat{A} = \sum_i \mathbf{Pr}(A = a_i) \times a_i$$

Where $\hat{A}$ is the predicted adjacency matrix, and a_i is the i-th network in the support of the model. This calculation is done for each individual network used in the model.

Value

A list of adjacency matrix of length nnets(object) or, if specified nnets(newdata).

Examples

```
data(fivenets)

# bernoulli graph
fit <- ergmito(fivenets ~ edges)

# all ties have the same likelihood
# which is roughly equal to:
# mean(nedges(fivenets)/(nvertex(fivenets)*(nvertex(fivenets) - 1)))
predict(fit)


# If we take into account vertex attributes, now the story is different!
fit <- ergmito(fivenets ~ edges + nodematch("female"))

# Not all ties have the same likelihood, since it depends on homophily!
predict(fit)
```

---

rbernoulli	*Random Bernoulli graph*

Description

Random Bernoulli graph

Usage

```
rbernoulli(n, p = 0.5)
```

Arguments

n	Integer vector. Size of the graph. If length(n) > 1, then it will a list of random graphs.
p	Probability of a tie. This may be either a scalar, or a vector of the same length of n.

Value

If n is a single number, a square matrix of size n with zeros in the diagonal. Otherwise it returns a list of length(n) square matrices of sizes equal to those specified in n.

Examples

```
# A graph of size 4
rbernoulli(4)

# 3 graphs of various sizes
rbernoulli(c(3, 4, 2))
```

```
# 3 graphs of various sizes and different probabilities
rbernoulli(c(3, 4, 6), c(.1, .2, .3))
```

---

same_dist   *Compare pairs of networks to see if those came from the same distri-*
                             bution

Description

If two networks are of the same size, and their vertex attributes are equal in terms of set comparison,
then we say those came from the same distribution

Usage

```
same_dist(net0, net1, attrnames = NULL, ...)
```

Arguments

net0, net1	Networks to be compared.
attrnames	Character vector. (optional) Names of the vertex attributes to be be compared on. This is ignored in the matrix case.
...	Ignored.

Details

This function is used during the call of [ergmito_formulae](#) to check whether the function can recycle
previously computed statistics for the likelihood function. In the case of models that only contain
structural terms, i.e. attribute less, this can save significant amount of computing time and memory.

Value

A logical with an attribute what. TRUE meaning that the two networks come from the same distribu-
tion, and FALSE meaning that they do not. If FALSE the what attribute will be equal to either "size"
or the name of the attribute that failed the comparison.

Examples

```
data(fivenets)
same_dist(fivenets[[1]], fivenets[[2]]) # Yes, same size
same_dist(fivenets[[1]], fivenets[[2]], "female") # No, different attr dist
```

---

simulate.ergmito   *Draw samples from a fitted* ergmito *model*

Description

Draw samples from a fitted ergmito model

Usage

```
## S3 method for class 'ergmito'
simulate(object, nsim = 1, seed = NULL, which_networks = 1L, theta = NULL, ...)
```

Arguments

object	An object of class ergmito.
nsim	Integer scalar. Number of samples to draw from the selected set of networks.
seed	See stats::simulate
which_networks	Integer vector. Specifies what networks to sample from. It must be within 1 and nnets(object).
theta, . . .	Further arguments passed to new_rergmito.

Examples

```
data(fivenets)
fit <- ergmito(fivenets ~ edges + nodematch("female"))

# Drawing 200 samples from networks 1 and 3 from the model
ans <- simulate(fit, nsim = 200, which_networks = c(1, 3))
```

---

vcov.ergmito   *Estimation of ERGMs using Maximum Likelihood Estimation (MLE)*

Description

ergmito uses Maximum Likelihood Estimation (MLE) to fit Exponential Random Graph Models
for single or multiple small networks, the later using pooled-data MLE. To do so we use exact like-
lihoods, which implies fully enumerating the support of the model. Overall, the exact likelihood
calculation is only possible when dealing with directed (undirected) networks size 5 (7). In gen-
eral, directed (undirected) graphs with more than 5 (7) vertices should not be fitted using MLE, but
instead other methods such as the MC-MLE algorithm or the Robbins-Monro Stochastic Approx-
imation algorithm, both of which are available in the ergm R package. The workhorse function of
ergmito is the ergm package function ergm::ergm.allstats().

Usage

```
## S3 method for class 'ergmito'
vcov(object, solver = NULL, ...)

ergmito(
  model,
  model_update = NULL,
  stats_weights = NULL,
  stats_statmat = NULL,
  optim.args = list(),
  init = NULL,
  use.grad = TRUE,
  target_stats = NULL,
  ntries = 1L,
  keep.stats = TRUE,
  target_offset = NULL,
  stats_offset = NULL,
  ...
)
```

Arguments

object	An object of class ergmito
solver	Function. Used to compute the inverse of the hessian matrix. When not null, the variance-covariance matrix is recomputed using that function. By default, ergmito uses MASS::ginv.
...	Further arguments passed to the method. In the case of ergmito, ... are passed to ergmito_formulae.
model	Model to estimate. See ergm::ergm. The only difference with ergm is that the LHS can be a list of networks.
model_update	A formula. this can be used to apply transformations, create interaction effects, add offset terms, etc. (see examples below and more details in ergmito_formulae).
stats_weights	Either an integer vector or a list of integer vectors (see exact_loglik).
stats_statmat	Either a matrix or a list of matrices (see exact_loglik).
optim.args	List. Passed to stats::optim.
init	A numeric vector. Sets the starting parameters for the optimization routine. Default is a vector of zeros.
use.grad	Logical. When TRUE passes the gradient function to optim. This is intended for testing only (internal use).
target_stats	A matrix of target statistics (see ergm::ergm).
ntries	Integer scalar. Number of tries to estimate the MLE (see details).
keep.stats	Logical scalar. When TRUE (the default), the matrices and vectors associated with the sufficient statistics will be returned. Otherwise the function discards them. This may be useful for saving memory space when estimating multiple models.

```
target_offset, stats_offset
                    See exact_loglik().
```

Details

The support of the sufficient statistics is calculated using ERGM's `ergm::ergm.allstats()` function.

Value

An list of class ergmito:

- `call` The program call.
- `coef` Named vector. Parameter estimates.
- `iterations` Integer. Number of times the log-likelihood was evaluated (see stats::optim).
- `mle.lik` Numeric. Final value of the objective function.
- `null.lik` Numeric. Final value of the objective function for the null model.
- `covar` Square matrix of size length(coef). Variance-covariance matrix computed using the exact hessian as implemented in exact_hessian.
- `coef.init` Named vector of length length(coef). Initial set of parameters used in the optimization.
- `formulae` An object of class ergmito_loglik.
- `nobs` Integer scalar. Number of networks in the model.
- `network` Networks passed via model.
- `optim.out,optim.args` Results from the optim call and arguments passed to it.
- `status,note` Convergence code. See check_convergence
- `best_try` Integer scalar. Index of the run with the highest log-likelihood value.
- `history` Matrix of size ntries * (k + 1). History of the parameter estimates and the reached log-likelihood values.
- `timer` Vector of times (for benchmarking). Each unit marks the starting point of the step.

Methods base::print(), base::summary(), stats::coef(), stats::logLik(), stats::nobs(), stats::vcov(), stats::AIC(), stats::BIC(), stats::confint(), and stats::formula() are available.

MLE

Maximum Likelihood Estimates are obtained using the stats::optim function. The default method for maximization is BFGS using both the log-likelihood function and its corresponding gradient.

Another important factor to consider is the existence of the MLE estimates As shown in Handcock (2003), if the observed statistics are near the border if the support function (e.g. too many edges or almost none), then, even if the MLE estimates exists, the optimization function may not be able to reach the optima. Moreover, if the target (observed) statistics live in the boundary, then the MLE estimates do not exists. In general, this should not be an issue in the context of the pooled model, as the variability of observed statistics should be enough to avoid those situations.

The function `ergmito` will try to identify possible cases of non-existence, of the MLE, and if identified then try to re estimate the model parameters using larger values than the ones obtained, if the log-likelihood is greater, then it is assumed that the model is degenerate and the corresponding values will be replaced with either +Inf or -Inf. By default, this behavior is checked anytime that the absolute value of the estimates is greater than 5, or the sufficient statistics were flagged as potentially outside of the interior of the support (close to zero or to its max).

In the case of `ntries`, the optimization is repeated that number of times, each time perturbing the `init` parameter by adding a Normally distributed vector. The result which reaches the highest log-likelihood will be the one reported as parameter estimates. This feature is intended for testing only. Anecdotally, `optim` reaches the max in the first try.

See Also

The function plot.ergmito() and gof_ergmito() for post-estimation diagnostics.

Examples

```
# Generating a small graph
set.seed(12)
n <- 4
net <- rbernoulli(n, p = .3)

model <- net ~ edges + mutual

library(ergm)
ans_ergmito <- ergmito(model)
ans_ergm   <- ergm(model)

# The ergmito should have a larger value
ergm.exact(ans_ergmito$coef, model)
ergm.exact(ans_ergm$coef, model)

summary(ans_ergmito)
summary(ans_ergm)

# Example 2: Estimating an ERGMito using data with know DGP parameters -----
data(fivenets)

model1 <- ergmito(fivenets ~ edges + nodematch("female"))
summary(model1) # This data has know parameters equal to -2.0 and 2.0

# Example 3: Likelihood ratio test using the lmtest R package

if (require(lmtest)) {
  data(fivenets)
  model1 <- ergmito(fivenets ~ edges + nodematch("female"))
  model2 <- ergmito(fivenets ~ edges + nodematch("female") + mutual)

  lrtest(model1, model2)
  # Likelihood ratio test
  #
```

```
  # Model 1: fivenets ~ edges + nodematch("female")
  # Model 2: fivenets ~ edges + nodematch("female") + mutual
  #   #Df  LogLik Df  Chisq Pr(>Chisq)
  # 1   2 -34.671
  # 2   3 -34.205 1 0.9312     0.3346
}

# Example 4: Adding an reference term for edge-count ---------------------

# Simulating networks of different sizes
set.seed(12344)
nets <- rbernoulli(c(rep(4, 10), rep(5, 10)), c(rep(.2, 10), rep(.1, 10)))

# Fitting an ergmito under the Bernoulli model
ans0 <- ergmito(nets ~ edges)
summary(ans0)
#
# ERGMito estimates
#
# formula:
#   nets ~ edges
#
#        Estimate Std. Error z value  Pr(>|z|)
# edges -1.68640    0.15396 -10.954 < 2.2e-16 ***
# ---
# Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
# AIC: 279.3753    BIC: 283.1436    (Smaller is better.)


# Fitting the model including a reference term for networks of size 5.
# Notice that the variable -n- and other graph attributes can be used
# with -model_update-.
ans1 <- ergmito(nets ~ edges, model_update = ~ I(edges * (n == 5)))
summary(ans1)
#
# ERGMito estimates
#
# formula:
#   nets ~ edges + I(edges * (n == 5))
#
#                     Estimate Std. Error z value  Pr(>|z|)
# edges               -1.18958    0.21583 -5.5116 3.556e-08 ***
# I(edges * (n == 5)) -0.90116    0.31250 -2.8837   0.00393 **
# ---
# Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
# AIC: 272.9916    BIC: 280.5282    (Smaller is better.)

# The resulting parameter for the edge-count is smaller for networks
# of size five
plogis(coef(ans1)[1])   # 0.23
plogis(sum(coef(ans1))) # 0.11

# We can see that in this case the difference in edge-count matters.
```

```
if (require(lmtest)) {

  lrtest(ans0, ans1)
  # Likelihood ratio test
  #
  # Model 1: nets ~ edges
  # Model 2: nets ~ edges + I(edges * (n == 5))
  # #Df  LogLik Df  Chisq Pr(>Chisq)
  # 1    1 -138.69
  # 2    2 -134.50  1 8.3837   0.003786 **
  #   ---
  #   Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
}
```

Index