

Introduction to Exponential-family Random Graph (ERG or p^*) modeling with *ergm*

Version 3.11.0

The Statnet Development Team

October 14, 2020

Contents

1	Getting Started	1
2	Statistical network modeling; the <i>ergm</i> command and <i>ergm</i> object	2
3	Model terms available for <i>ergm</i> estimation and simulation	10
3.1	Terms provided with <i>ergm</i>	10
3.2	Coding new terms	11
4	Network simulation: the <i>simulate</i> command and <i>network.list</i> objects	11
5	Examining the quality of model fit – <i>GOF</i>	13
6	Diagnostics: troubleshooting and checking for model degeneracy	17
7	Working with egocentrically sampled network data	25
8	Additional functionality in the statnet family of packages	26
8.1	In the <i>ergm</i> and <i>network</i> packages:	26
8.2	In other packages from the statnet suite:	26
8.2.1	Static (cross-sectional) network analysis packages	26
8.2.2	Dynamic (longitudinal) network analysis packages:	27
8.3	R packages that build on statnet	27
8.4	Additional functionality in development:	27
9	Statnet Commons: The core development team	27

1 Getting Started

This vignette is based on the *ergm* tutorial presented at INSNA Sunbelt - St. Pete Beach, Florida, Feb 2011.

Open an R session, and set your working directory to the location where you would like to save this work. You can do this with the pull-down menus (File¿Change Dir) or with the command:

```
setwd('full.path.for.the.folder')
```

To install all of the packages in the statnet suite:

```
install.packages('statnet')  
library(statnet)
```

Or, to only install the specific statnet packages needed for this tutorial:

```
install.packages('network')  
install.packages('ergm')  
install.packages('sna')  
library(network)  
library(ergm)  
library(sna)
```

After the first time, to update the packages one can either repeat the commands above, or use:

```
update.packages('name.of.package')
```

For this tutorial, we will need one additional package (coda), which is recommended (but not required) by ergm:

```
install.packages('coda')  
library(coda)
```

2 Statistical network modeling; the *ergm* command and *ergm* object

Make sure the statnet package is attached:

```
library(statnet)
```

or


```
library(ergm)  
library(sna)  
  
## Loading required package: statnet.common  
##  
## Attaching package: 'statnet.common'  
## The following object is masked from 'package:base':  
##  
## order  
## sna: Tools for Social Network Analysis  
## Version 2.6 created on 2020-10-5.  
## copyright (c) 2005, Carter T. Butts, University of California-Irvine  
## For citation information, type citation("sna").  
## Type help(package="sna") to get started.  
  
set.seed(1)
```

The ergm package contains several network data sets that you can use for practice examples.

```
data(package='ergm') # tells us the datasets in our packages
data(florentine) # loads flomarriage and flobusiness data
flomarriage # Let's look at the flomarriage data

## Network attributes:
##   vertices = 16
##   directed = FALSE
##   hyper = FALSE
##   loops = FALSE
##   multiple = FALSE
##   bipartite = FALSE
##   total edges= 20
##     missing edges= 0
##     non-missing edges= 20
##
## Vertex attribute names:
##   priorates totalties vertex.names wealth
##
## No edge attributes

plot(flomarriage) # Let's view the flomarriage network
```



Remember the general ergm representation of the probability of the observed network, and the conditional log-odds of a tie:

$$\Pr(Y = y) = \exp[\theta'g(y)]/k(\theta)$$

Y is a network; $g(y)$ is a vector of network stats; θ is the vector of coefficients; $k(\theta)$ is a normalizing constant.

$$\text{logit}(\Pr(Y_{ij} = 1|Y^c)) = \theta' \Delta(g(y))_{ij}$$

Y_{ij} is an actor pair in Y ; Y^c is the rest of the network; $\Delta(g(y))_{ij}$ is the change in $g(y)$ when the value of Y_{ij} is toggled on.

We begin with the simplest possible model, the Bernoulli or Erdős-Rényi model, which contains only an edge term.

```
flomodel.01 <- ergm(flomarriage~edges) # fit model

## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
```

```

## Maximizing the pseudolikelihood.
## Finished MPLE.
## Stopping at the initial estimate.
## Evaluating log-likelihood at the estimate.

flomodel.01

##
## Call:
## ergm(formula = flomarriage ~ edges)
##
##
## MLE Coefficients:
## edges
## -1.609

summary(flomodel.01) # look in more depth

## Call:
## ergm(formula = flomarriage ~ edges)
##
## Iterations: 5 out of 20
##
## Monte Carlo MLE Results:
##      Estimate Std. Error MCMC % z value Pr(>|z|)
## edges  -1.6094    0.2449      0  -6.571  <1e-04 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##      Null Deviance: 166.4  on 120  degrees of freedom
##      Residual Deviance: 108.1  on 119  degrees of freedom
##
## AIC: 110.1    BIC: 112.9    (Smaller is better.)

```

How to interpret this model? The log-odds of any tie occurring is:

$$\begin{aligned}
 & -1.609 \times \text{change in the number of ties} \\
 & = -1.609 \times 1
 \end{aligned}$$

for all ties, since the addition of any tie to the network changes the number of ties by 1!

Corresponding probability is:

$$\begin{aligned}
 & \exp(-1.609)/(1 + \exp(-1.609)) \\
 & = 0.1667
 \end{aligned}$$

which is what you would expect, since there are 20/120 ties.

Let's add a term often thought to be a measure of "clustering": the number of completed triangles. Note we're in stochastic simulation now – your output will differ

```

flomodel.02 <- ergm(flomarriage~edges+triangle)

## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
## Maximizing the pseudolikelihood.
## Finished MPLE.
## Starting Monte Carlo maximum likelihood estimation (MCMLE):
## Iteration 1 of at most 20:
## Optimizing with step length 1.
## The log-likelihood improved by 0.002182.
## Step length converged once. Increasing MCMC sample size.
## Iteration 2 of at most 20:
## Optimizing with step length 1.
## The log-likelihood improved by 0.0008553.
## Step length converged twice. Stopping.
## Finished MCMLE.
## Evaluating log-likelihood at the estimate. Using 20 bridges: 1 2 3 4 5 6 7 8 9 10
11 12 13 14 15 16 17 18 19 20 .
## This model was fit using MCMC. To examine model diagnostics and
## check for degeneracy, use the mcmc.diagnostics() function.

summary(flomodel.02)

## Call:
## ergm(formula = flomarriage ~ edges + triangle)
##
## Iterations: 2 out of 20
##
## Monte Carlo MLE Results:
##           Estimate Std. Error MCMC % z value Pr(>|z|)
## edges      -1.6774     0.3458     0  -4.851  <1e-04 ***
## triangle    0.1623     0.5608     0   0.289    0.772
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Null Deviance: 166.4 on 120 degrees of freedom
## Residual Deviance: 108.1 on 118 degrees of freedom
##
## AIC: 112.1 BIC: 117.6 (Smaller is better.)

coef1 = flomodel.02$coef[1]
coef2 = flomodel.02$coef[2]
logodds = coef1 + c(0,1,2) * coef2
expit = function(x) 1/(1+exp(-x))
ps = expit(logodds)
coef1 = round(coef1, 3)
coef2 = round(coef2, 3)

```

```
logodds = round(logodds, 3)
ps = round(ps, 3)
```

Again, how to interpret coefficients?

Conditional log-odds of two actors forming a tie is:

$$-1.677 \times \text{change in the number of ties} + 0.162 \times \text{change in number of triangles}$$

- if the tie will not add any triangles to the network, its log-odds is: -1.677 .
- if it will add one triangle to the network, its log-odds is: $-1.677 + 0.162 = -1.515$
- if it will add two triangles to the network, its log-odds is: $-1.677 + 0.162 \times 2 = -1.353$
- the corresponding probabilities are 0.157, 0.18, and 0.205.

Let's take a closer look at the ergm object itself:

```
class(flomodel.02) # this has the class ergm

## [1] "ergm"

names(flomodel.02) # let's look straight at the ERGM obj.

## [1] "coef"          "sample"        "sample.obs"    "iterations"
## [5] "MCMCtheta"     "loglikelihood" "gradient"       "hessian"
## [9] "covar"         "failure"       "network"       "newnetworks"
## [13] "newnetwork"    "coef.init"     "est.cov"       "coef.hist"
## [17] "stats.hist"    "steplen.hist"  "control"       "etamap"
## [21] "call"         "ergm_version"  "MPLE_is_MLE"   "formula"
## [25] "target.stats"  "target.esteq"  "constrained"   "constraints"
## [29] "reference"     "estimate"      "offset"        "drop"
## [33] "estimable"     "null.lik"      "mle.lik"

flomodel.02$coef

##      edges  triangle
## -1.6774030  0.1623367

flomodel.02$formula

## flomarriage ~ edges + triangle


flomodel.02$mle.lik

## 'log Lik.' -54.03713 (df=2)

wealth <- flomarriage %v% 'wealth' # the %v% extracts vertex
wealth # attributes from a network

## [1] 10 36 55 44 20 32 8 42 103 48 49 3 27 10 146 48

plot(flomarriage, vertex.cex=wealth/25) # network plot with vertex size
```



proportional to wealth

We can test whether edge probabilities are a function of wealth:

```
flomodel.03 <- ergm(flomarriage~edges+nodecov('wealth'))

## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
## Maximizing the pseudolikelihood.
## Finished MPLE.
## Stopping at the initial estimate.
## Evaluating log-likelihood at the estimate.

summary(flomodel.03)

## Call:
## ergm(formula = flomarriage ~ edges + nodecov("wealth"))
##
```



```
## Iterations: 4 out of 20
##
## Monte Carlo MLE Results:
##           Estimate Std. Error MCMC % z value Pr(>|z|)
## edges      -2.594929   0.536056      0  -4.841  <1e-04 ***
## nodecov.wealth 0.010546   0.004674      0   2.256   0.0241 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Null Deviance: 166.4 on 120 degrees of freedom
## Residual Deviance: 103.1 on 118 degrees of freedom
##
## AIC: 107.1 BIC: 112.7 (Smaller is better.)
```

Yes, there is a significant positive wealth effect on the probability of a tie.

Let's try a model or two on:

Is there a statistically significant tendency for ties to be reciprocated ('mutuality')?

```
data(samplk)
ls() # directed data: Sampson's Monks
samplk3
plot(samplk3)
sampmodel.01 <- ergm(samplk3~edges+mutual)
summary(sampmodel.01)
```

Let's try a larger network

```
data(faux.mesa.high)
mesa <- faux.mesa.high
```

```
plot(mesa)
mesa
plot(mesa, vertex.col='Grade')
legend('bottomleft',fill=7:12,legend=paste('Grade',7:12),cex=0.75)
fauxmodel.01 <- ergm(mesa ~edges + nodematch('Grade',diff=T) + nodematch('Race',diff=T))
summary(fauxmodel.01)
```

Note that two of the coefficients are estimated as -Inf (the nodematch coefficients for race Black and Other). Why is this?

```
table(mesa %v% 'Race') # Frequencies of race

##
## Black  Hisp NatAm Other White
##      6   109    68     4    18

mixingmatrix(mesa, "Race")
```

```
## Note: Marginal totals can be misleading
## for undirected mixing matrices.
##      Black Hisp NatAm Other White
## Black      0      8     13      0      5
## Hisp       8     53     41      1     22
## NatAm      13     41     46      0     10
## Other       0      1      0      0      0
## White       5     22     10      0      4
```

So the problem is that there are very few students in the Black and Other race categories, and these students form no homophilous (within-group) ties. The empty cells are what produce the -Inf estimates.

Time to consider some missing data:

```
missnet <- network.initialize(10,directed=F)
missnet[1,2] <- missnet[2,7] <- missnet[3,6] <- 1
missnet[4,6] <- missnet[4,9] <- NA
missnet
plot(missnet)
ergm(missnet~edges)
```

The coefficient equals -2.590. This is the log-odds of the probability .0698. Our network has 3 ties, out of the 43 nodal pairs (10 choose 2 minus 2) whose dyad status we have observed. $3/43 = 0.0698$.

```
ergm(missnet~edges+degree(2))
missnet[4,6] <- missnet[4,9] <- 0
ergm(missnet~edges+degree(2))
```

The two estimates for the degree 2 coefficient differ considerably. In the first case, there is one node we know for sure has degree 2, two that may or may not, and seven that we know for sure do not. In the latter, there is one node that has degree 2, and nine that do not.

3 Model terms available for *ergm* estimation and simulation

Model terms are the expressions (e.g. “triangle”) used to represent predictors on the right-hand side of equations used in:

- calls to **ergm** (to estimate an ergm model)
- calls to **simulate** (to simulate networks from an ergm model fit)
- calls to **summary** (to obtain measurements of network statistics on a dataset)

3.1 Terms provided with *ergm*

For a list of available terms that can be used to specify an ERGM, see Appendix B, or type:

```
help('ergm-terms')
```

For a more complete discussion of these terms see the ‘Specifications’ paper in J Stat Software v. 24. (link is available online at www.statnet.org)

3.2 Coding new terms

The package (`ergm.userterms`) and tutorial are aimed at making it much easier than before to write one's own terms. This package is available on CRAN, and installing it will also download the tutorial (`ergmuserterms.pdf`). We teach a workshop at the Sunbelt meetings. Note that writing up new `ergm` terms requires some knowledge of C and the ability to build R from source (although the latter is covered in the tutorial).

4 Network simulation: the *simulate* command and *network.list* objects

Once we have estimated the coefficients of an ERGM, the model is completely specified. It defines a probability distribution across all networks of this size. If the model is a good fit to the observed data, then networks drawn from this distribution will be more likely to "resemble" the observed data. To see examples of networks drawn from this distribution we use the `simulate` command:

```
flomodel.03.sim <- simulate(flomodel.03,nsim=10)
class(flomodel.03.sim)

## [1] "network.list"

summary(flomodel.03.sim)

## Number of Networks: 10
## Model: flomarriage ~ edges + nodecov("wealth")
## Reference: ~Bernoulli
## Constraints: ~.
## Parameters:
##      edges nodecov.wealth
## -2.59492903    0.01054591
##
## Stored network statistics:
##      edges nodecov.wealth
## [1,]    11          1110
## [2,]    16          1690
## [3,]    24          2174
## [4,]    21          1840
## [5,]    21          2334
## [6,]    16          2143
## [7,]    19          2085
## [8,]    20          1878
## [9,]    22          1798
## [10,]   18          2084
## attr(,"monitored")
## [1] FALSE FALSE
## Number of Networks: 10
## Model: flomarriage ~ edges + nodecov("wealth")
## Reference: ~Bernoulli
## Constraints: ~.
## Parameters:
```

```

##          edges nodecov.wealth
##    -2.59492903    0.01054591

length(flomodel.03.sim)

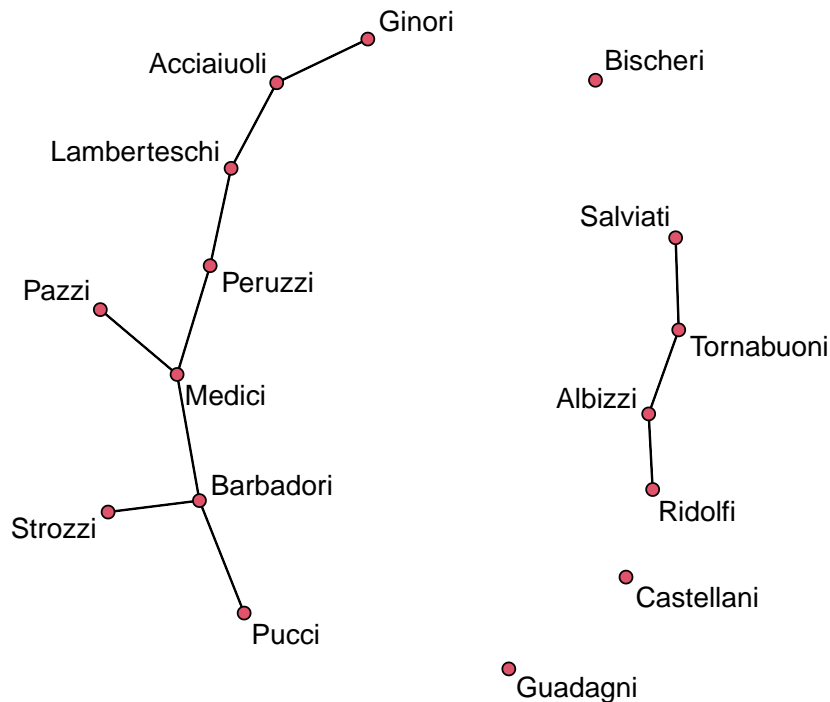
## [1] 10

flomodel.03.sim[[1]]

## Network attributes:
##   vertices = 16
##   directed = FALSE
##   hyper = FALSE
##   loops = FALSE
##   multiple = FALSE
##   bipartite = FALSE
##   total edges= 11
##   missing edges= 0
##   non-missing edges= 11
##
## Vertex attribute names:
##   priorates totalties vertex.names wealth
##
## No edge attributes

plot(flomodel.03.sim[[1]], label= flomodel.03.sim[[1]] %v% "vertex.names")

```



Voila. Of course, yours will look somewhat different.

5 Examining the quality of model fit – *GOF*

ERGMs are generative models – that is, they represent the process that governs tie formation at a local level. These local processes in turn aggregate up to produce characteristic global network properties, even though these global properties are not explicit terms in the model. One test of whether a model “fits the data” is therefore how well it reproduces these global properties. We do this by choosing a network statistic that is not in the model, and comparing the value of this statistic observed in the original network to the distribution of values we get in simulated networks from our model.

```

flomodel.03.gof <- gof(flomodel.03~degree)


flomodel.03.gof

##
## Goodness-of-fit for degree
##
##   obs min mean max MC p-value


```

```
## 0 1 0 1.17 5 1.00
## 1 4 0 3.65 9 1.00
## 2 2 0 4.10 10 0.34
## 3 6 0 3.05 6 0.06
## 4 2 0 2.12 10 1.00
## 5 0 0 1.02 4 0.72
## 6 1 0 0.53 2 0.88
## 7 0 0 0.20 2 1.00
## 8 0 0 0.12 1 1.00
## 9 0 0 0.02 1 1.00
## 10 0 0 0.02 1 1.00
##
## Goodness-of-fit for model statistics
##
## obs min mean max MC p-value
## edges 20 12 20.25 33 1.00
## nodecov.wealth 2168 1104 2182.19 3700 0.98
```

```
plot(flomodel.03.gof)
```




Goodness-of-fit diagnostics




```
mesamodel.02 <- ergm(mesa~edges)

## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
## Maximizing the pseudolikelihood.
## Finished MPLE.
## Stopping at the initial estimate.
## Evaluating log-likelihood at the estimate.

mesamodel.02.gof <- gof(mesamodel.02~distance, control=control.gof.ergm(nsim=10))
plot(mesamodel.02.gof)
```



Goodness-of-fit diagnostics



For a good example of model exploration and fitting for the Add Health Friendship networks, see Goodreau, Kitts & Morris, *Demography* 2009.

6 Diagnostics: troubleshooting and checking for model degeneracy

The computational algorithms in `ergm` use MCMC to estimate the likelihood function. Part of this process involves simulating a set of networks to approximate unknown components of the likelihood.

When a model is not a good representation of the observed network the estimation process may be affected. In the worst case scenario, the simulated networks will be so different from the observed network that the algorithm fails altogether. This can occur for two general reasons. First, the simulation algorithm may fail to converge, and the sampled networks are thus not from the specified distribution. Second, the model parameters used to simulate the networks are too different from the MLE, so even though the simulation algorithm is producing a representative sample of networks, this is not the sample that would be produced under the MLE.

For more detailed discussions of model degeneracy in the ERGM context, see the papers in *J Stat Software* v. 24. (link is available online at www.statnet.org)

We can use diagnostics to see what is happening with the simulation algorithm, and these can lead us to ways to improve it.

We will first consider a simulation where the algorithm works. To understand the algorithm, consider

```
fit <- ergm(flobusiness~edges+degree(1),  
  control=control.ergm(MCMC.interval=1, MCMC.burnin=1000, seed=1))
```

This runs a version with every network returned. Let us look at the diagnostics produced:


```
mcmc.diagnostics(fit, center=F)
```

Let's look more carefully at a default model fit:

```
fit <- ergm(flobusiness~edges+degree(1))  
mcmc.diagnostics(fit, center=F)
```

Now let us look at a more interesting case, using a larger network:

```
data('faux.magnolia.high')  
magnolia <- faux.magnolia.high  
plot(magnolia, vertex.cex=.5)
```



```
fit <- ergm(magnolia~edges+triangle, control=control.ergm(seed=1))

## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
## Maximizing the pseudolikelihood.
## Finished MPLE.
## Starting Monte Carlo maximum likelihood estimation (MCMLE):
## Iteration 1 of at most 20:
## Error in ergm.MCMLE(init, nw, model, initialfit = (initialfit <- NULL), : Number of
edges in a simulated network exceeds that in the observed by a factor of more than 20.
This is a strong indicator of model degeneracy or a very poor starting parameter configuration.
If you are reasonably certain that neither of these is the case, increase the MCMLE.density.guard
control.ergm() parameter.
```

Very interesting. This model produced degenerate networks. You could have gotten some more feedback about this during the fitting, by using:

```
fit <- ergm(magnolia~edges+triangle, control=control.ergm(seed=1), verbose=T)
```

You might try to increase the MCMC sample size:

```
fit <- ergm(magnolia~edges+triangle,seed=1,
  control = control.ergm(seed=1, MCMC.samplesize=20000),
  verbose=T)
mcmc.diagnostics(fit, center=F)
```

How about trying the more robust version of modeling triangles: GWESP? (For a technical introduction to GWESP see Hunter and Handcock; for a more intuitive description and empirical application see Goodreau Kitts and Morris 2009)

```
fit <- ergm(magnolia~edges+gwap(0.5,fixed=T),
  control = control.ergm(seed=1))
mcmc.diagnostics(fit)
```

Still degenerate, but maybe getting closer?

```
fit <- ergm(magnolia~edges+gwap(0.5,fixed=T)+nodematch('Grade')+nodematch('Race')+
  nodematch('Sex'),
  control = control.ergm(seed=1),
  verbose=T)

pdf('diagnostics1.pdf') #Use the recording function if possible, otherwise send to pdf
mcmc.diagnostics(fit)
dev.off() #If you saved to pdf, look at the file

fit <- ergm(magnolia~edges+gwap(0.25,fixed=T)+nodematch('Grade')+nodematch('Race')+
  nodematch('Sex'),
  control = control.ergm(seed=1))
mcmc.diagnostics(fit)
```

One more try...

```
fit <- ergm(magnolia~edges+gwap(0.25,fixed=T)+nodematch('Grade')+nodematch('Race')+
  nodematch('Sex'),
  control = control.ergm(seed=1,MCMC.samplesize=4096,MCMC.interval=8192),
  verbose=T)

## Evaluating network in model.
## Initializing Metropolis-Hastings proposal(s): ergm:MH.TNT
## Initializing model.
## Using initial method 'MPLE'.
## Fitting initial model.
## Starting maximum pseudolikelihood estimation (MPLE):
## Evaluating the predictor and response matrix.
## Maximizing the pseudolikelihood.
## Finished MPLE.
## Starting Monte Carlo maximum likelihood estimation (MCMLE):
```

```

## Density guard set to 19563 from an initial count of 974 edges.
##
## Iteration 1 of at most 20 with parameter:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      -9.6939224             1.6624034         2.8170436         0.9870588
##      nodematch.Sex
##      0.6188166
## Starting unconstrained MCMC...
## Back from unconstrained MCMC.
## Average estimating function values:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      25.60278             -41.99772         20.71875         27.33154
##      nodematch.Sex
##      -24.64038
## Starting MCML E Optimization...
## Optimizing with step length 0.800609232668285.
## Using lognormal metric (see control.ergm function).
## Using log-normal approx (no optim)
## The log-likelihood improved by 4.953.
##
## Iteration 2 of at most 20 with parameter:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      -9.7584644             1.7754153         2.7683026         0.9152198
##      nodematch.Sex
##      0.7406027
## Starting unconstrained MCMC...
## Back from unconstrained MCMC.
## Average estimating function values:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      27.11841             11.96378         25.95801         22.29419
##      nodematch.Sex
##      10.88306
## Starting MCML E Optimization...
## Optimizing with step length 1.
## Using lognormal metric (see control.ergm function).
## Using log-normal approx (no optim)
## The log-likelihood improved by 0.3807.
## Step length converged once. Increasing MCMC sample size.
##
## Iteration 3 of at most 20 with parameter:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      -9.7838929             1.7923641         2.7481880         0.9135958
##      nodematch.Sex
##      0.7666666
## Starting unconstrained MCMC...
## Back from unconstrained MCMC.
## Average estimating function values:
##      edges gwesp.fixed.0.25  nodematch.Grade  nodematch.Race
##      -8.553772             -9.831627         -8.513489         -7.437805
##      nodematch.Sex
##      -5.701660

```

```

## Starting MCMLE Optimization...
## Optimizing with step length 1.
## Using lognormal metric (see control.ergm function).
## Using log-normal approx (no optim)
## Starting MCMC s.e. computation.
## The log-likelihood improved by 0.02911.
## Step length converged twice. Stopping.
## Finished MCMLE.
## Evaluating log-likelihood at the estimate. Using 20 bridges: 1 2 3 4 5 6 7 8 9 10
11 12 13 14 15 16 17 18 19 20 .
## This model was fit using MCMC. To examine model diagnostics and
## check for degeneracy, use the mcmc.diagnostics() function.

```

```
mcmc.diagnostics(fit)
```

```

## Sample statistics summary:
##
## Iterations = 131072:134340608
## Thinning interval = 8192
## Number of chains = 1
## Sample size per chain = 16384
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##              Mean      SD Naive SE Time-series SE
## edges          -8.554 48.00   0.3750         4.331
## gwesp.fixed.0.25 -9.832 43.52   0.3400         4.523
## nodematch.Grade -8.513 45.92   0.3587         4.403
## nodematch.Race  -7.438 44.41   0.3470         4.199
## nodematch.Sex   -5.702 39.42   0.3080         3.711
##
## 2. Quantiles for each variable:
##
##              2.5%    25%    50%    75% 97.5%
## edges          -102.00 -41.00 -8.00 24.00 85.00
## gwesp.fixed.0.25 -91.40 -39.44 -9.83 19.01 80.81
## nodematch.Grade -97.42 -39.00 -8.00 22.00 82.00
## nodematch.Race  -92.00 -38.00 -8.00 22.00 81.00
## nodematch.Sex   -82.00 -33.00 -7.00 21.00 72.00
##
##
## Sample statistics cross-correlations:
##              edges gwesp.fixed.0.25 nodematch.Grade nodematch.Race
## edges          1.0000000          0.8573291          0.9631962          0.9464149
## gwesp.fixed.0.25 0.8573291          1.0000000          0.8775269          0.8550061
## nodematch.Grade 0.9631962          0.8775269          1.0000000          0.9210282
## nodematch.Race  0.9464149          0.8550061          0.9210282          1.0000000
## nodematch.Sex   0.9197596          0.8127159          0.8903887          0.8738815
##
##              nodematch.Sex
## edges          0.9197596

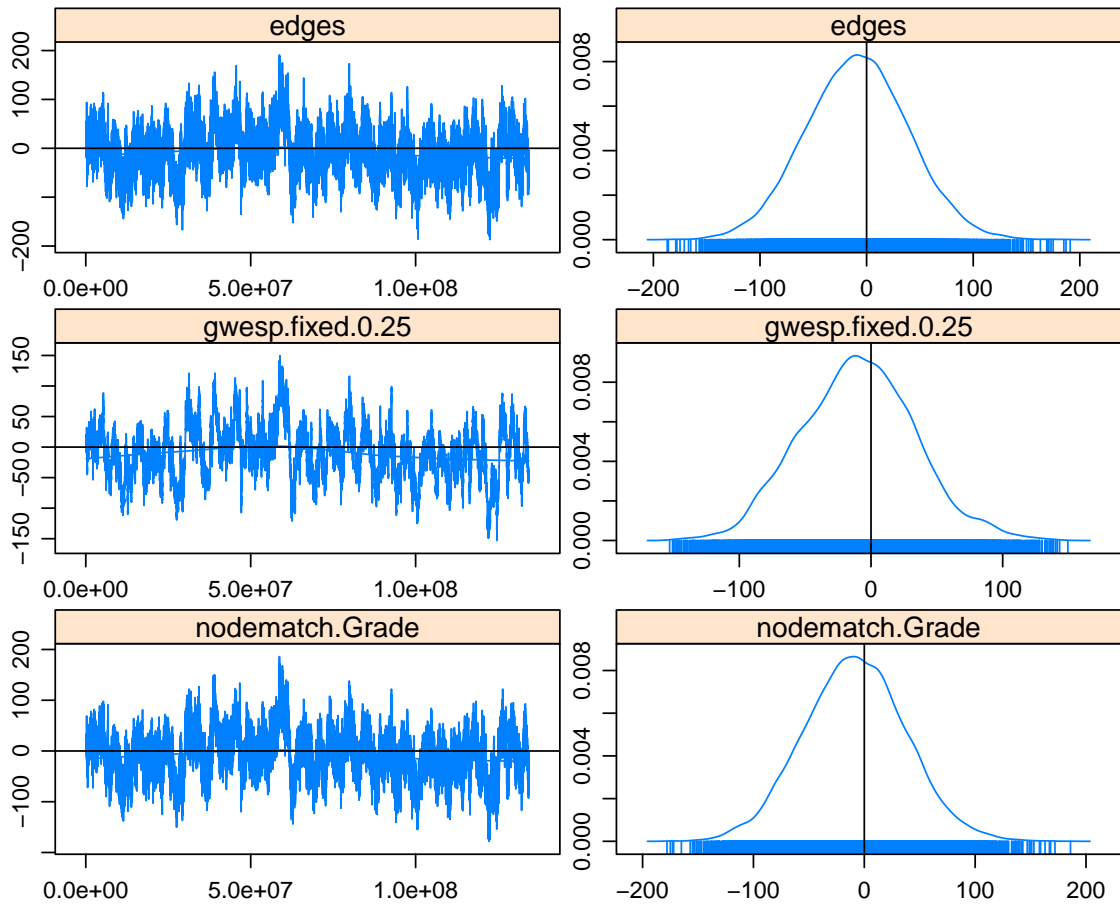
```

```

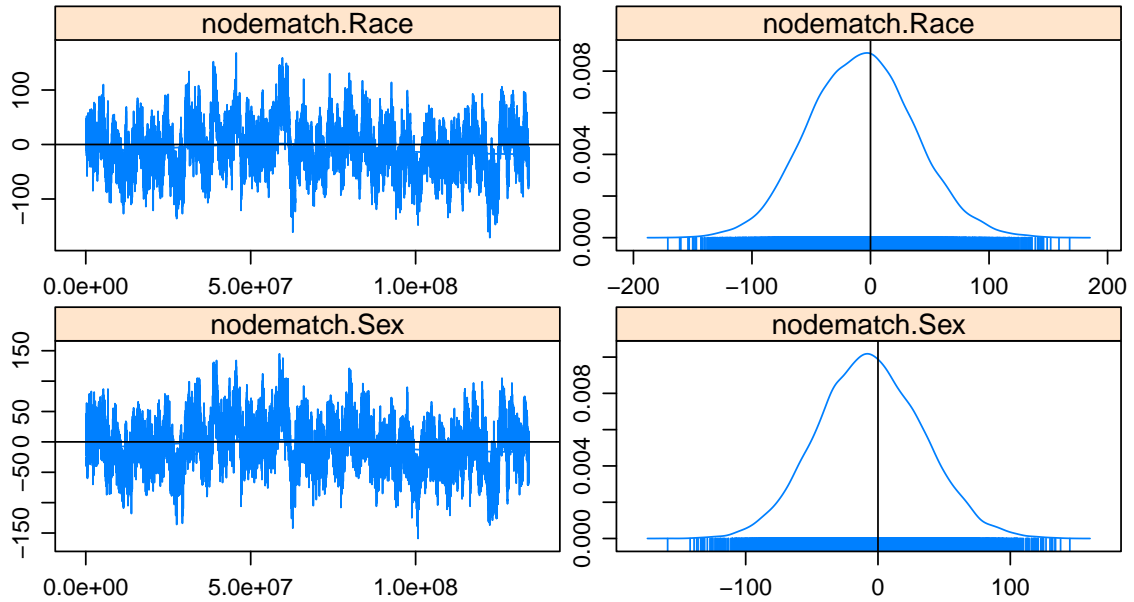
## gwesp.fixed.0.25      0.8127159
## nodematch.Grade      0.8903887
## nodematch.Race       0.8738815
## nodematch.Sex        1.0000000
##
## Sample statistics auto-correlation:
## Chain 1
##          edges gwesp.fixed.0.25 nodematch.Grade nodematch.Race
## Lag 0      1.0000000      1.0000000      1.0000000      1.0000000
## Lag 8192   0.7894063      0.9831165      0.8402751      0.8278664
## Lag 16384  0.7410728      0.9679183      0.7928438      0.7797591
## Lag 24576  0.7198520      0.9535880      0.7718407      0.7590152
## Lag 32768  0.7039626      0.9402569      0.7570195      0.7438275
## Lag 40960  0.6935912      0.9275207      0.7472533      0.7320818
##          nodematch.Sex
## Lag 0      1.0000000
## Lag 8192   0.8025758
## Lag 16384  0.7522586
## Lag 24576  0.7276672
## Lag 32768  0.7133930
## Lag 40960  0.7001329
##
## Sample statistics burn-in diagnostic (Geweke):
## Chain 1
##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
##          edges gwesp.fixed.0.25 nodematch.Grade nodematch.Race
##          0.2207      0.2015      0.3108      0.3655
##          nodematch.Sex
##          0.2931
##
## Individual P-values (lower = worse):
##          edges gwesp.fixed.0.25 nodematch.Grade nodematch.Race
##          0.8253354      0.8403417      0.7559479      0.7147684
##          nodematch.Sex
##          0.7694554
## Joint P-value (lower = worse): 0.4691064 .

```

Sample statistics



Sample statistics



```
##
```

```
## MCMC diagnostics shown here are from the last round of simulation, prior to computation of final p
```

Success! Of course, in real life one might have a lot more trial and error.

Changes in version 3.2 of the `ergm` estimation algorithm mean that the MCMC diagnostic plots can no longer be used to ensure that the mean statistics from the model match the observed network statistics. For that functionality, please use the GOF command: `gof(fit, GOF= model)`. The plots can still be used to assess MCMC mixing and convergence.

7 Working with egocentrically sampled network data

One of the most powerful features of ERGMs is that they can be used to estimate models from from egocentrically sampled data, and the fitted models can then be used to simulate complete networks (of any size) that will have the properties of the original network that are observed and represented in the model.

In many empirical contexts, it is not feasible to collect a network census or even an adaptive (link-traced) sample. Even when one of these may be possible in practice, egocentrically sampled data are typically cheaper and easier to collect.

Long regarded as the poor country cousin in the network data family, egocentric data contain a remarkable amount of information. With the right statistical methods, such data can be used to explore the properties of the complete networks in which they are embedded. The basic idea here is to combine what is observed, with assumptions, to define a class of models that describe the distribution of networks that are centered on the observed properties. The variation in these networks quantifies some of the uncertainty introduced by the assumptions.

The egocentric estimation/simulation framework extends to temporal ERGMs (“TERGMs”) as well, with the minimal addition of an estimate of partnership duration. This makes it possible to simulate complete dynamic networks from a single cross-sectional egocentrically sampled network. For an example of what you can do with this, check out the network movie we developed to explore the impact of dynamic network structure on HIV transmission, see <https://statnet.org/movies>.

While the `ergm` package has had this capability for many years (and old ERGM workshops had a section on this), there is now a specific package that makes this much easier: ‘`ergm.ego`’. The new package includes accurate statistical inference (so you can get standard errors for model coefficients), and many utilities that simplify the task of reading in the data, conducting exploratory analyses, and specifying model options.

We now have a separate workshop/tutorial for ‘`ergm.ego`’, so we no longer cover this material in the current ERGM workshop. As always, this workshop material can be found online at the `statnet` wiki.

8 Additional functionality in the `statnet` family of packages

8.1 In the `ergm` and `network` packages:

- ERGMs for valued ties – see the paper by Pavel Krivitsky (2012)
- Analysis of bipartite networks – `network` recognizes this as an attribute of the network and `ergm` provides specific model terms for such networks that begin with `b1` or `b2`
(try: `search.ergmTerms(categories=c('bipartite'))`).

8.2 In other packages from the `statnet` suite:

These are in stand-alone packages that can be downloaded either from CRAN, or from the `statnet` website. Many have online training materials from our workshops. For more detailed information, please visit the `statnet` webpage (<https://statnet.org>).

8.2.1 Static (cross-sectional) network analysis packages

- `sna` – Traditional SNA methods and summaries.
- `latentnet` – Latent space and latent cluster analysis.
- `netperm` – Network permutation models.
- `degreenet` – MLE estimation for degree distributions (negative binomial, Poisson, scale-free, etc.)
- `networksis` – Simulation of bipartite networks with given degree distributions.

8.2.2 Dynamic (longitudinal) network analysis packages:

- **tergm** – Temporal ERGMs (TERGMs): discrete-time dynamic network models for longitudinal network panel data, and other temporal extensions.
- **relevent** – Relational event models: continuous-time dynamic network models for longitudinal network data.
- **networkDynamic** – Dynamic network data storage and manipulation.
- **ndtv** – Network movie maker.

8.3 R packages that build on statnet

There is a growing number of R packages written by other folks that build on or extend the functionality of the statnet suite. You can get a current list of those packages by looking at the reverse depends/suggests on CRAN. A partial list includes:

- **EpiModel** package – Mathematical Modeling of Infectious Disease, includes functions for deterministic compartmental modeling, stochastic individual contact modeling, and stochastic network modeling
- **RDS** package – Estimation with data collected using Respondent-Driven Sampling.
- **Bergm** package – Bayesian ERGM estimation
- **hergm** package – Hierarchical Exponential-Family Random Graph Models with Local Dependence (for latent groups).
- **lvm4net** package – Latent variable models.
- **VBLPCM** package – Variational Bayes Latent Position Cluster Models.
- **xergm** package – Temporal exponential random graph models (TERGM) by bootstrapped pseudolikelihood, MCMC MLE and (temporal) network autocorrelation models.

8.4 Additional functionality in development:

- **ergm.ego** package – ERGM estimation and inference from egocentrically sampled data (expected May 2015)
- **tsna** package – Temporally extended (vertex and edge) SNA methods for dynamic longitudinal network data (expected May 2015)
- **MLergm** package – ERGM estimation and inference for multi-level data (for observed groups) (expected 2016)

9 Statnet Commons: The core development team

Mark S. Handcock jhandcock@stat.ucla.edu
David R. Hunter jdhunter@stat.psu.edu
Carter T. Butts jbuttsc@uci.edu
Steven M. Goodreau jgoodreau@u.washington.edu
Skye Bender-deMoll jskyebend@skyeome.net
Martina Morris jmorris@u.washington.edu
Pavel N. Krivitsky jpavel@uow.edu.au

Appendix A: Clarifying the terms – ergm and network

You will see the terms `ergm` and `network` used in multiple contexts throughout the documentation. This is common in R, but often confusing to newcomers. To clarify:

`ergm`

- **ERGM**: the acronym for an Exponential Random Graph Model; a statistical model for relational data that takes a generalized exponential family form.
- **ergm package**: one of the packages within the `statnet` suite
- **ergm function**: a function within the `ergm` package; fits an ERGM to a network object, creating an `ergm` object in the process.
- **ergm object**: a class of objects produced by a call to the `ergm` function, representing the results of an ERGM fit to a network.

`network`

- **network**: a set of actors and the relations among them. Used interchangeably with the term graph.
- **network package**: one of the packages within the `statnet` suite; used to create, store, modify and plot the information found in network objects.
- **network object**: a class of object in R used to represent a network.

References

- Goodreau, S., J. Kitts and M. Morris (2009). Birds of a Feather, or Friend of a Friend? Using Statistical Network Analysis to Investigate Adolescent Social Networks. *Demography* 46(1):103-125.
- Handcock, M. S., D. R. Hunter, C. T. Butts, S. M. Goodreau and M. Morris (2008). `statnet`: Software Tools for the Representation, Visualization, Analysis and Simulation of Network Data. *Journal of Statistical Software* 42(01).
- Hunter DR, Handcock MS, Butts CT, Goodreau SM, Morris M (2008b). `ergm`: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks. *Journal of Statistical Software*, 24(3). <https://www.jstatsoft.org/v24/i03/>.
- Krivitsky, P.N.(2009). PhD thesis. *University of Washington, Seattle, WA*
- Krivitsky, P. N., M. S. Handcock and M. Morris (2011). Network Size and Composition Effects in Exponential-Family Random Graph Models. *Statistical Methodology* 8:319-339
- Krivitsky PN (2012). Exponential-family random graph models for valued networks. *Electronic Journal of Statistics* 6:1100-1128