An AMEN tutorial (based on notes by P. Hoff)

Alexander Volfovsky

March 2019

AME model

The data

- There are n actors/nodes labeled $1, \ldots, n$
- Y is a sociomatrix: y_{ij} is a dyadic relationship between node i and node j.
- y_{ii} frequently undefined.
- Covariates:
  - node specific: x_i
  - dyad specific: x_{ij}

Social relations model

- Goal: describe the variability in Y.
- Sender effects describe sociability.
- Receiver effects describe popularity.

Capture this in the Social Relations Model (SRM)

$$y_{ij} = a_i + b_j + \epsilon_{ij}$$

Almost an ANOVA — want to relate a_i to b_i since the senders/receivers are from the same set.

Social relations model

$$
\begin{aligned}
y_{ij} =& \mu + a_i + b_j + \epsilon_{ij} \\
(a_i, b_i) \overset{iid}{\sim}& N(0, \Sigma_{ab}) \\
(\epsilon_{ij}, \epsilon_{ji}) \overset{iid}{\sim}& N(0, \Sigma_e)
\end{aligned}
$$

- $\Sigma_{ab} = \begin{pmatrix} \sigma_a^2 & \sigma_{ab} \\ \sigma_{ab} & \sigma_b^2 \end{pmatrix}$ describes sender/receiver variability and within person similarity.
- $\Sigma_e = \sigma_\epsilon^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ describes within dyad correlation.

Variability

$$
\begin{aligned}
var(y_{ij}) =& \sigma_a^2 + \sigma_b^2 + \sigma_\epsilon^2 \\
cov(y_{ij}, y_{ik}) =& \sigma_a^2 \\
cov(y_{ij}, u_{kj}) =& \sigma_b^2 \\
cov(y_{ij}, y_{jk}) =& \sigma_{ab} \\
cov(y_{ij}, y_{ji}) =& 2\sigma_{ab} + \rho\sigma_\epsilon^2
\end{aligned}
$$

How hard is it to fit this model?

```
fit_SRM <- ame(Y)
```

Pictures that pop up

These help capture how well the Markov Chain is mixing and goodness of fit information.

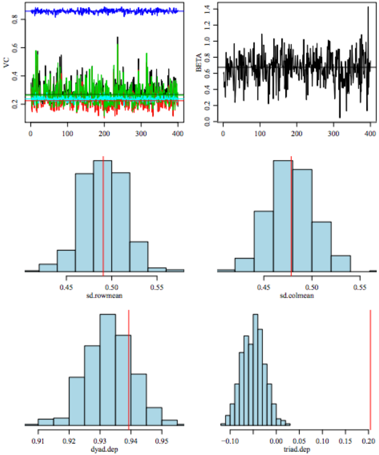

Figure 2: Default plots generated by the ame command.

Goodness of fit

Posterior predictive distributions.

- sd.rowmean: standard deviation of row means of Y.
- sd.colmean: standard deviation of column means of Y.
- dyad.dep: correlation between vectorized Y and vectorized Y^t
- triad.dep:

$$\frac{\sum_i \sum_{jk} e_{ij} e_{jk} e_{ki}}{\#\text{triangle on } n \text{ nodes}} Var(vec(Y))^{3/2}$$

Incorporating covariates

Imagine you have some covariates and want to fit

$$y_{ij} = \beta_d^t x_{d,ij} + \beta_r^t x_{r,i} + \beta_c^t x_{c,j} + a_i + b_j + \epsilon_{ij}$$

- $x_{d,ij}$ are dyad specific covariates.
- $x_{r,i}$ are row (sender) covariates.
- $x_{c,i}$ are column (receiver) covariates.
- Frequently $x_{r,i} = x_{c,i} = x_i$
- When does this not make sense?
- (Example: popularity is affected by athletic success, but sociability is not)

How hard is it to fit this model?

```
fit_SRRM <- ame(Y, Xd=Xd,Xr=Xr,Xc=Xc)
```

Parsing the input

```
fit_SRRM <- ame(Y,
Xdyad=Xd, #n x n x pd array of covariates
Xrow=Xr,  #n x pr matrix of nodal row covariates
Xcol=Xc   #n x pc matrix of nodal column covariates
)
```

- $Xr_{i,p}$ is the value of the pth row covariate for node i.
- $Xd_{i,j,p}$ is the value of the pth dyadic covariate in the direction of i to j.

Back to basics

Can you get rid of the dependencies in the model?

```
    fit_rm<-ame(Y,Xd=Xd,Xr=Xn,Xc=Xn,
rvar=FALSE, #should you fit row random effects?
cvar=FALSE, #should you fit column random effects?
dcor=FALSE  #should you fit a dyadic correlation?
)
```

Note that summary will output:

```
    Variance parameters:
pmean psd
va 0.000 0.000
cab 0.000 0.000
vb 0.000 0.000
rho 0.000 0.000
ve 0.229 0.011
```

Introducing multiplicative effects

- SR(R)M can represent second-order dependencies very well.
- Has a hard time capturing "triadic" behavior.
- Homophily: create dyadic covariates $x_{d,ij} = x_i x_j$
- Generally this can be represented by $x_{r_i}^t B x_{j,i} = \sum_k \sum_l b_{kl} x_{r,ik} x_{c,jl}$
- This is linear in the covariates and so can be baked into the amen framework.
- Sometimes there is excess correlation to account.
- This suggests a multiplicative effects model:

$$y_{ij} = \beta_d^t x_{d,ij} + \beta_r^t x_{r,i} + \beta_c^t x_{c,j} + a_i + b_j + u_i^t v_j + \epsilon_{ij}$$

Fitting these models and beyond

```
fit_ame2<-ame(Y,Xd,Xn,Xn,
R=2 #dimension of the multiplicative effect
)
```

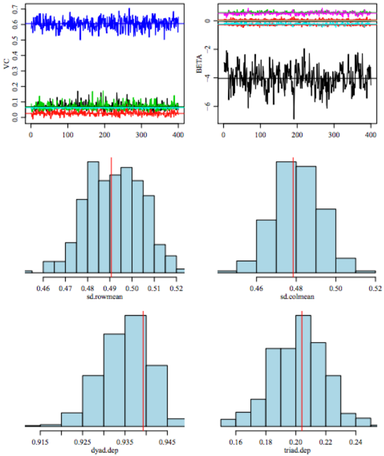

Ordinal models

Imagine a binary (probit) model:

$$y_{ij} = 1_{z_{ij}>0} \quad z_{ij} = \mu + a_i + b_j + \epsilon_{ij}$$

Looks like the SRM on the latent scale.

```
fit_SRM<-ame(Y,
model="bin" #lots of model options here
)
```

If we go to the iid set up this is just an Erdos-Renyi model:

```
fit_SRG<-ame(Y,model="bin",
rvar=FALSE,cvar=FALSE,dcor=FALSE)
```