

Using multivariate autoregressive
models to estimate interaction
strengths, aka the B matrix
Part II: community dynamics

Eli Holmes

FISH 507 – Applied Time Series Analysis

28 February 2017

Topics

- MAR(1) models as sum of community interactions
- MAR(1) models as a general approximation to nonlinear systems
- Stability properties of MAR(1) models
- Estimating MAR(1) models

The theory in this lecture draws largely from

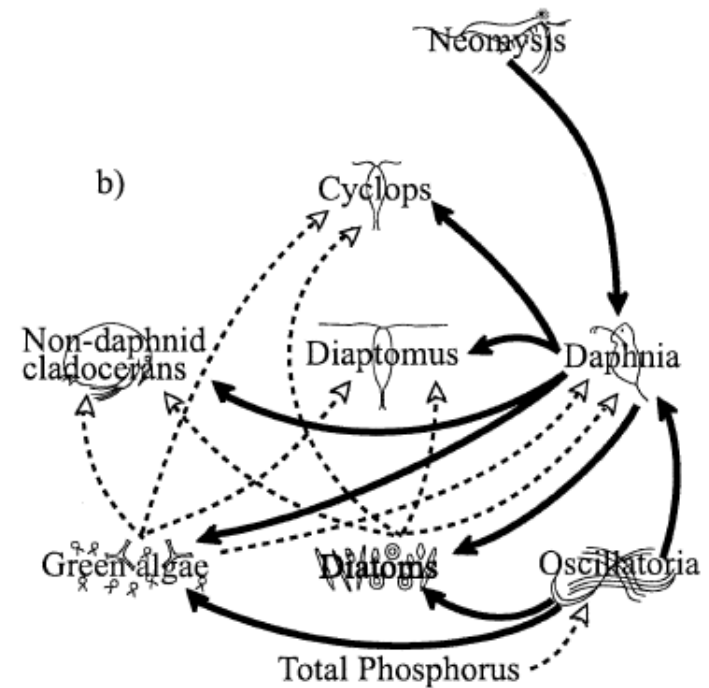
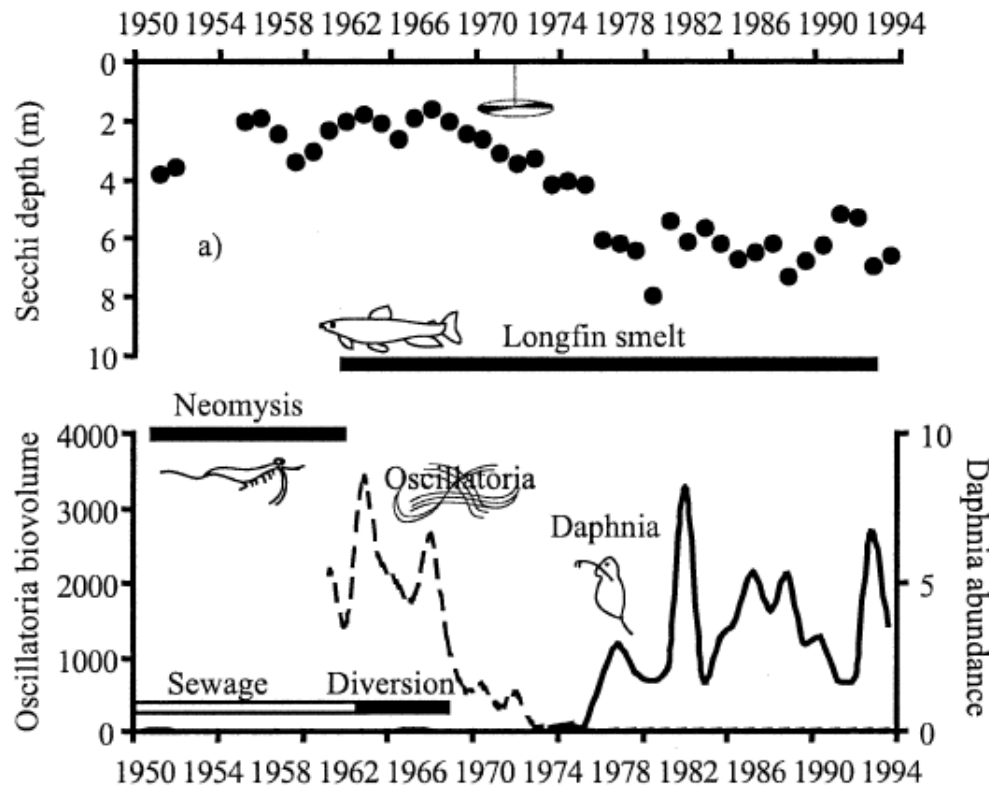
Ives AR, Dennis B, Cottingham KL, Carpenter SR. 2003. Estimating community stability and ecological interactions from time series data. *Ecological Monographs* 73: 301-330

Many ecological applications are reviewed in

Hampton, S.E., E.E. Holmes, D.E. Pendleton, L.P. Scheef, M.D. Scheuerell, and E.J. Ward. 2013. Quantifying effects of abiotic and biotic drivers on community dynamics with multivariate autoregressive (MAR) models. *Ecology* 94:2663-2669


Links to both are on the course syllabus

Lake Washington: a large change in sewage inputs in the late 1960s led to a dramatic change in the plankton community



Hampton et al. (2006)

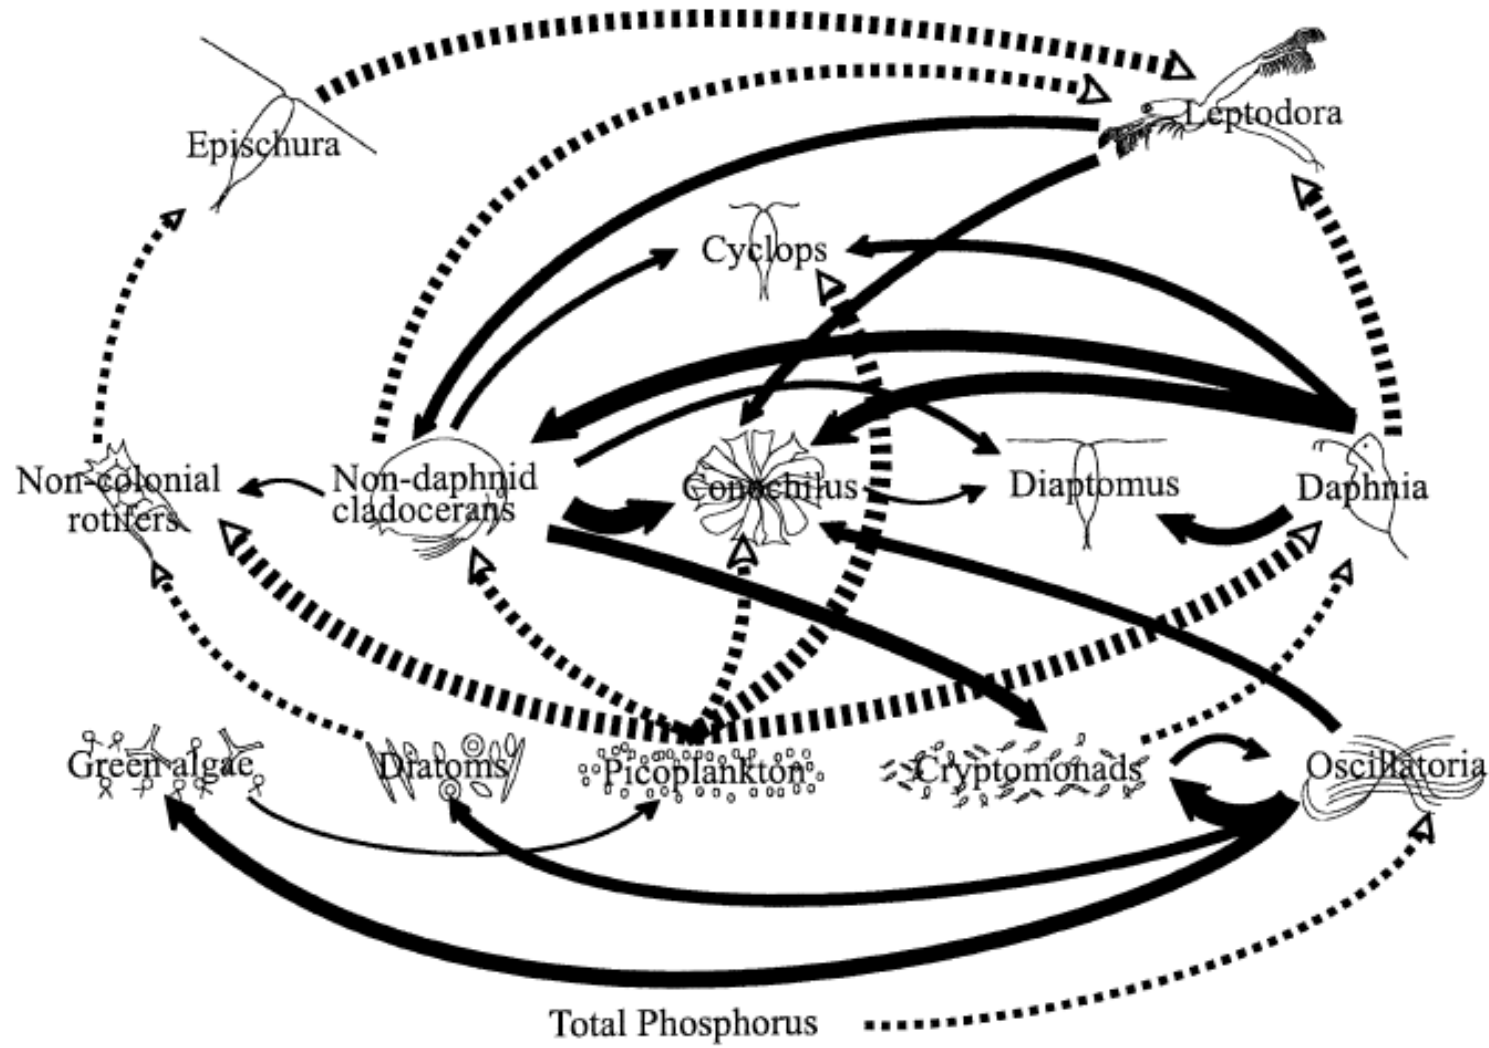
How can we make sense of this?



Dashed lines are (+) interactions; solid lines are (-)

Hampton et al. (2006)

And reduce to the 'important' links?



Dashed lines are (+) interactions; solid lines are (-)

Hampton et al. (2006)

Multivariate AR(1) process, “MAR1”

Note no SS part

Assume p different interacting species, then for each sp. i :

$$\begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_n \end{bmatrix}_t = \begin{bmatrix} b_{11} & b_{21} & \dots & b_{n1} \\ b_{12} & b_{22} & \dots & b_{n2} \\ \dots & \dots & \dots & \dots \\ b_{1n} & b_{2n} & \dots & b_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_n \end{bmatrix}_{t-1} + \begin{bmatrix} u_1 \\ u_2 \\ \dots \\ u_n \end{bmatrix} + \begin{bmatrix} w_1 \\ w_2 \\ \dots \\ w_n \end{bmatrix}_t$$

B

$$\mathbf{w}_t \sim MVN(0, \mathbf{Q})$$

$X_{j,t-1}$ is abundance of species j at time $t-1$

b_{ij} is effect of species j on species i

Note: when $i = j$, b_{ij} is effect of species i on itself

Multivariate AR(1) process

Can add effects of q different environmental covariates (U):

$$\begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_n \end{bmatrix}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u} + \begin{bmatrix} c_{11} & c_{21} \\ c_{12} & c_{22} \\ \dots & \dots \\ c_{1n} & c_{2n} \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}_t + \begin{bmatrix} w_1 \\ w_2 \\ \dots \\ w_n \end{bmatrix}_t$$

$q=2$

$c_{k,t}$ is covariate k (e.g., temperature) at time t

C_{ik} is effect of covariate k on species i

NOTE: covariate not time lagged; if you want a time-lagged covariate you need to use the time-lagged covariate as c_t

Multivariate AR(1) process

Can write all of this as:

$$\mathbf{x}_t = \mathbf{u} + \mathbf{B}\mathbf{x}_{t-1} + \mathbf{C}\mathbf{c}_t + \mathbf{w}_t \quad \mathbf{w}_t \sim MVN(\mathbf{0}, \mathbf{Q})$$

\mathbf{x}_t is $p \times 1$ vector of species abundance at time t

\mathbf{u} is $p \times 1$ vector of intrinsic growth rates

\mathbf{B} is $p \times p$ matrix of density-dependent effects

\mathbf{x}_{t-1} is $p \times 1$ vector of species abundance at time $t-1$

\mathbf{C} is $p \times q$ matrix of covariate effects

\mathbf{c}_t is $q \times 1$ vector of covariate values at time t

Multivariate AR(1) process

A note on the B matrix structure:

$$\mathbf{B} = \begin{bmatrix} b_{11} & b_{12} & b_{13} & \cdots & b_{1p} \\ b_{21} & b_{22} & \cdots & \cdots & \vdots \\ b_{31} & \vdots & b_{33} & \cdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b_{p1} & \cdots & \cdots & \cdots & b_{pp} \end{bmatrix}$$


Inter-specific effects
(can be set to zero)

Inter-specific effects
(can be set to zero)

Intra-specific effects

Real systems do not necessarily have log-linear density-dependence

Note: If N is not spanning very large values a linear model may suffice



Modern literature on MAR(1) models allow B to be time-varying which allows that linear approximation to vary in time.

Stationary distribution for MAR(1)

Mean vector: $\mathbf{m}_\infty = (\mathbf{I} - \mathbf{B})^{-1} \mathbf{u}$

Covariance matrix: $\mathbf{V}_\infty = \mathbf{B}\mathbf{V}_\infty\mathbf{B}' + \mathbf{Q}$

Note: cannot solve analytically this for \mathbf{V}_∞

Solution: use the $\text{Vec}(\bullet)$ operator

$$\text{Vec}(\mathbf{V}_\infty) = (\mathbf{I} - \mathbf{B} \otimes \mathbf{B})^{-1} \text{Vec}(\mathbf{Q})$$

(analogous to the univariate case)

Stability properties of MAR(1) models

- Our interest is in stable systems (i.e. all eigenvalues of \mathbf{B} lie within unit circle)
- Stability can be measured in several ways, but here are three that we will use:

Stability measure	More stable when...
Variance	variance of the stationary distribution is low relative to that for the process error
Return rate	rapid approach to the stationary distribution (i.e. high return rate)
Reactivity	fewer departures from the mean of the stationary distribution (i.e. low reactivity)

Stability properties of MAR(1) models

Return rates (ie, resilience)

- In deterministic models, equilibrium is a point or stable limit cycle
- In stochastic models, equilibrium is the stationary distribution
- Rate of return to the stochastic equilibrium is the rate at which the transition distribution converges to the stationary distribution from an initial, known value
- The more rapid the convergence, the more stable the system




Fig. 2B - Ives et al. (2003)

Stability properties of MAR(1) models

Return rates

- Rate of return increases as $|b|$ gets smaller
- Rate of return to mean governed by dominant eigenvalue of $\mathbf{B} \equiv \max(\lambda_{\mathbf{B}})$
- Rate of return of covariance defined by dominant eigenvalue of $\mathbf{B} \otimes \mathbf{B} \equiv \max(\lambda_{\mathbf{B} \otimes \mathbf{B}})$

Take home msg: you can write the return rate as a function of B. So if we can estimate B, we can say something about the stability of the system as measure by return rate.





Fig. 3 - Ives et al. (2003)

Stability properties of MAR(1) models

Reactivity

A “reactive” system moves away from a stable equilibrium following a perturbation, even though the system will eventually return to equilibrium.

Low reactivity



High reactivity




Fig. 2C - Ives et al. (2003)

High reactivity occurs when species interactions greatly amplify the environmental forcing

Not how fast does it return but how far does it go?

Stability properties of MAR(1) models

Reactivity

Squared Euclidian distance for \mathbf{X}_{t-1} from μ_∞

$$= \|\mathbf{X}_{t-1} - \mu_\infty\|^2$$

Squared Euclidian distance for \mathbf{X}_t from μ_∞

$$= \|\mathbf{u} + \mathbf{B}\mathbf{X}_{t-1} - \mu_\infty\|^2 = \|\mathbf{B}(\mathbf{X}_{t-1} - \mu_\infty)\|^2$$

Change in squared Euclidian distance

$$= \|\mathbf{B}(\mathbf{X}_{t-1} - \mu_\infty)\|^2 - \|\mathbf{X}_{t-1} - \mu_\infty\|^2$$

Expectation of this difference

$$= -\frac{\text{tr}[\mathbf{Q}]}{\text{tr}[\mathbf{V}_\infty]} \leq \boxed{\max(\lambda_{\mathbf{B}'\mathbf{B}})} - 1$$

“worst-case”
scenario

Another form of reactivity is given in Neubert et al. (2009)




Fig. 3 - Ives et al. (2003)

Stability properties of MAR(1) models

Variance comparison

- Generally interested in the time spent away from equilibrium
- More time \rightarrow greater variance




Fig. 2A - Ives et al. (2003)

Combines properties of both return rate and reactivity

Stability properties of MAR(1) models

Variance comparison

1) $\tilde{v}_\infty = \frac{\psi^2}{(1-\lambda^2)}$ Smaller $\lambda \rightarrow$ smaller $v \rightarrow$ more stable

2) $\det(\mathbf{V}_\infty - \mathbf{Q}) = \det(\mathbf{V}_\infty) \det(\mathbf{B})^2$

$$\det(\mathbf{B})^2 = \frac{\det(\mathbf{V}_\infty - \mathbf{Q})}{\det(\mathbf{V}_\infty)} = (\lambda_1 \lambda_2 \lambda_3 \cdots \lambda_p)^2$$

To standardize across various systems

$$\det(\mathbf{B})^{2/p} = (\lambda_1 \lambda_2 \lambda_3 \cdots \lambda_p)^{2/p}$$

Smaller $\lambda \rightarrow$ $\det(\mathbf{B})^{2/p}$ \rightarrow more stable

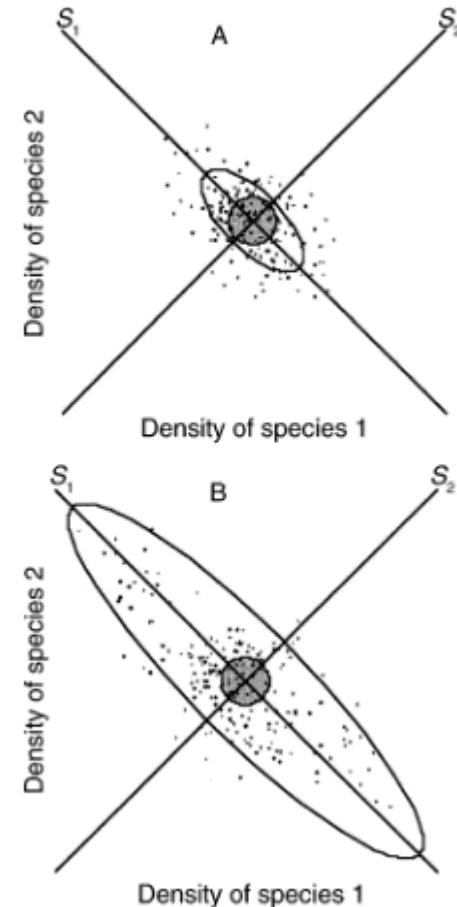


Fig. 3 - Ives et al. (2003)

MAR(1) parameter estimation

$$\mathbf{x}_t = \mathbf{u} + \mathbf{B}\mathbf{x}_{t-1} + \mathbf{C}\mathbf{c}_t + \mathbf{w}_t$$

Maximum likelihood estimation (MLE)

Goal: find the B and C, that maximize the likelihood conditioned on the observed data.

Pro: does not require the assumption of no observation error (we can add an observation process) and allows missing values.

Con: Slow.

MAR(1) parameter estimation

dat="some m x T matrix of log species abundances"

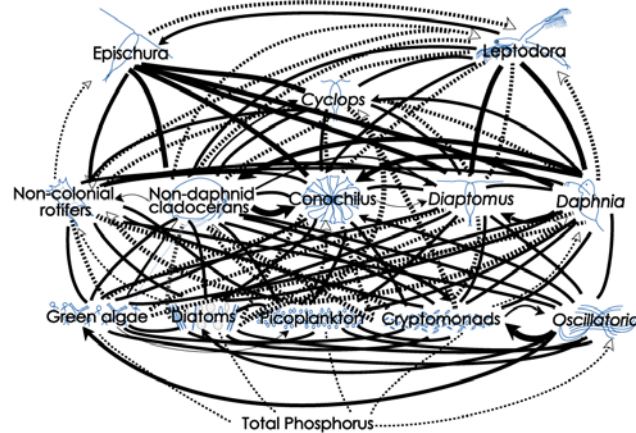
```
mod.list = list(  
  B="unconstrained",  
  U="zero",  
  Q="diagonal and unequal",  
  Z="identity",  
  A="zero",  
  R="zero",  
  C="unconstrained", c="some q x T matrix of covariates",  
  x0=dat[,1,drop=FALSE], tinitx=1  
)
```

MARSS(dat, model=mod.list)

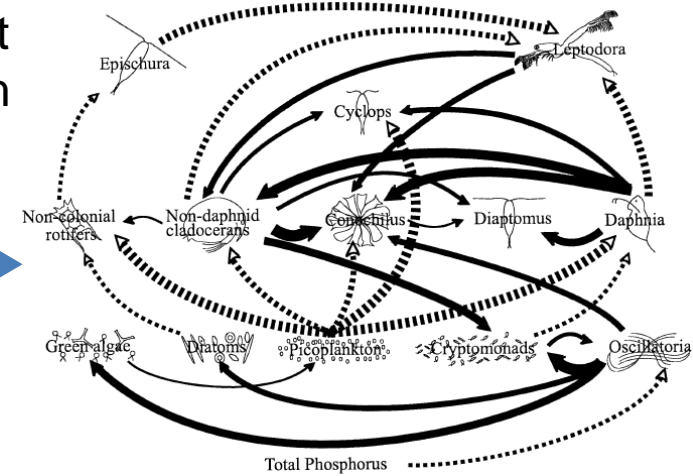
maximum-likelihood estimation via Expectation-
Maximization algorithm or method="BFGS" to use optim()
and a quasi-Newton method.

Goal is to find a simpler system to explain the observed species changes

13 x 13 B matrix



Compare model support of models with different B's set to zero



156 non-diagonal B terms. There are approximately $1e+47$ (that's 1 followed by 47 zeros) unique B matrices with different 0s.

- 1) We cannot do a brute force fitting of all models.
- 2) We should try to set some of those Bs to 0 via prior knowledge about the system
- 3) We need a fast estimation algorithm

Much faster MAR(1) parameter estimation

Conditional least squares (CLS)

The CLS parameter estimates are those values that minimize the squared difference between

- 1) the observed log population abundances at time t , and
- 2) those predicted by the MAR(1) model, conditional on the log population abundances at time $t-1$.

Requires the assumption of no observation error and no missing values allowed.

Conditional least squares (CLS)

$$\begin{array}{lll}
 \mathbf{Y}_2^T = [\mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_T] & m \times (T-1) & (m \times 1)(1 \times T-1) \\
 \mathbf{w} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_T] & m \times (T-1) & m \times q \\
 \mathbf{c} = [\mathbf{c}_2, \mathbf{c}_3, \dots, \mathbf{c}_T] & q \times (T-1) & m \times m
 \end{array}
 \quad \mathbf{Y}_2^T = \mathbf{u}\mathbf{1} + \mathbf{B}\mathbf{Y}_1^{T-1} + \mathbf{C}\mathbf{c} + \mathbf{w}$$

We work row by row through the equation and solve for each row of \mathbf{u} , \mathbf{B} and \mathbf{C} .

$$\text{Row } i: \text{ find } \min \left[\left(\mathbf{Y}_i - \hat{\mathbf{Y}}_i \right) \left(\mathbf{Y}_i - \hat{\mathbf{Y}}_i \right)^T \right] \text{ where } \hat{\mathbf{Y}}_i = \hat{u}_i + \hat{\mathbf{B}}_i \mathbf{X} + \hat{\mathbf{C}}_i \mathbf{c}$$

$$\text{Let } \mathbf{Z}^T = [1, \mathbf{Y}_1^{T-1}, \mathbf{c}] \text{ and } \hat{\mathbf{D}}_i = [\hat{u}_i, \hat{\mathbf{B}}_i, \hat{\mathbf{C}}_i] \text{ then } \hat{\mathbf{D}}_i = \mathbf{Y}_i \mathbf{Z}^T (\mathbf{Z}^T \mathbf{Z})^{-1}$$

Note Ives et al. 2003 writes this transposed.

MAR(1) parameter estimation

Conditional least squares (CLS)

There are a couple of advantages to solving this via CLS:

- 1) It's VERY fast, and
- 2) It's easy to zero out elements of B to “test” for their significance

$$\mathbf{B} = \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} \\ b_{21} & b_{22} & b_{23} & b_{24} \\ b_{31} & b_{32} & b_{33} & b_{34} \\ b_{41} & b_{42} & b_{43} & b_{44} \end{bmatrix}$$

MAR(1) parameter estimation

Conditional least squares (CLS)

There are a couple of advantages to solving this via CLS:

- 1) It's VERY fast, and
- 2) It's easy to zero out elements of B to “test” for their significance

$$\mathbf{B} = \begin{bmatrix} b_{11} & 0 & b_{13} & b_{14} \\ b_{21} & b_{22} & b_{23} & 0 \\ b_{31} & 0 & b_{33} & b_{34} \\ 0 & b_{42} & 0 & b_{44} \end{bmatrix}$$

MAR(1) parameter estimation

General steps to a MAR(1) analysis

- 1) Look at the data. Hopefully it looks like a mean-reverting process.
- 2) Deal with issues like seasonality and trends (look at papers to study how other have done this....)
- 3) Demean and standardize the variance of log population size and covariates
- 4) Impose biological constraints. “zooplankton are not affected by nitrogen”, “fish do not eat green algae”
- 5) Model selection for the B matrix

Search strategy for large model spaces

Best of 100 random search

- 1) 100 models are generated by randomly including (50/50) each element.
- 2) Best model (lowest AIC or BIC) is saved
- 3) Repeat 100 times to get a set of 100 “best” model.
- 4) Remove any variables that appears in less than 15 of the 100 best models.
- 5) Repeat the ‘Best of 100’ search (steps 1-3) with the new smaller set of variables.
- 6) Repeat steps 1-5 until number of variables in the 100 best is stable and chose best model.

Search strategy for large model spaces

Exhaustive via “leap and bound” algorithm

- 1) Algorithms to find best model over all possible variable combinations without actually doing exhaustive search.

Forward step-wise

- 1) Start with no off-diagonal b .
- 2) Add the b that most improves (reduces) model AIC
- 3) Stop at some step-AIC threshold

Backward step-wise

- 1) Same as forward but start with full model (all b 's)

Model diagnostics

Once selected, a MAR(1) model should be scrutinized

Useful diagnostics include:

- 1) Are residuals temporally autocorrelated (via ACF)?
- 2) Are mean or variance of the residuals correlated with any variates or covariates (via X-Y plots)?
- 3) Are residuals normally distributed (via normal probability plots)?
- 4) What proportion of the variance is explained by the model (via conditional R^2)?

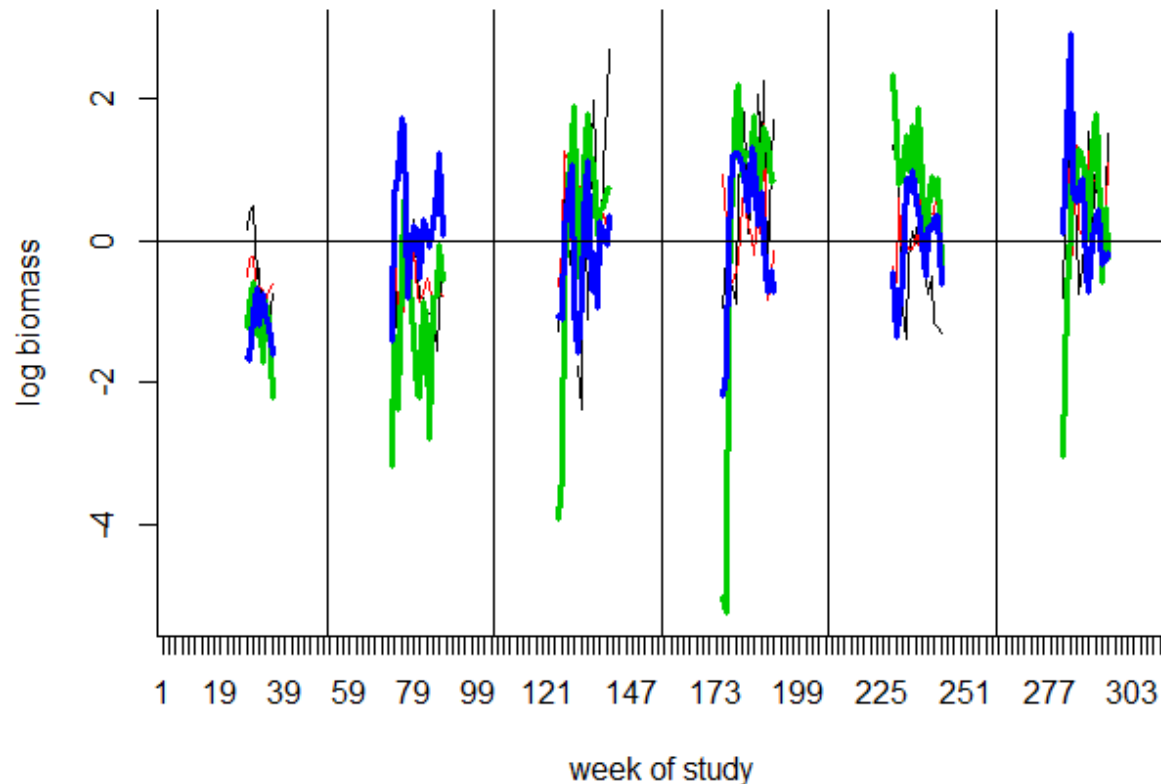
Note: conditional R^2 measures the proportion of change in log density from time $t-1$ to t

Model diagnostics

$$\mathbf{x}_t = \mathbf{u} + \mathbf{B}\mathbf{x}_{t-1} + \mathbf{C}\mathbf{c}_t + \mathbf{w}_t$$

Example of demeaned plankton data, 4 spp (Ives et al 2003)

- we are fitting one B matrix (and u)
- so level changes need to be explained by the covariates



Parameter confidence intervals (CI)

Can obtain CI's from model in 3 ways:

- 1) "Innovations" bootstrapping (how Ives et al (2003) does it)
 - Parameter estimation produces state residuals (w_t)
 - Randomly resample the residuals (by column) to produce a new \mathbf{X} (meaning use as your w_t)
 - Use new \mathbf{X} to again estimate parameters
 - Repeat "many" times
 - Estimate CI's from distribution of parameter estimates

Parameter confidence intervals (CI)

- 2) Parametric bootstrapping if you use a MLE method
 - Generate parametric bootstrap data from the MLE model
 - Use new data to again estimate parameters (bootstrap parameter estimates)
 - Repeat “many” times

Parameter confidence intervals (CI)

- 3) Estimate CI's from the (estimated) distribution of parameter estimates
 - Estimate the Hessian matrix (there are various R functions to do this)
 - Use that to compute the (estimated) CIs
 - Large- n approximation

MAR(1) with observation error

This is simply our (now) familiar MARSS model:

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{C}\mathbf{c}_t + \mathbf{w}_t, \text{ where } \mathbf{w}_t \sim \text{MVN}(\mathbf{0}, \mathbf{Q})$$

$$\mathbf{y}_t = \mathbf{x}_t + \mathbf{v}_t, \text{ where } \mathbf{v}_t \sim \text{MVN}(\mathbf{0}, \mathbf{R})$$

See the MARSS User Guide for a computer lab that walks through

- Fit community interaction model to freshwater plankton data from Ives et al. (2003)
- Add covariates to the above model
- Add an additional covariate (fish) observed with error
- Compute stability metrics
- MAR1 package on CRAN will apply the Ives et al 2003 algorithms to data sets.