Lecture 14
Bayesian Models for Spatio-Temporal Data

Dennis Sun
Stats 253

August 13, 2014

Outline of Lecture

1. Recap of Bayesian Models
2. Empirical Bayes
3. Case 1: Long-Lead Forecasting of Sea Surface Temperatures
4. Case 2: Modeling and Forecasting the Eurasian Dove Invasion
5. Case 3: Mediterranean Surface Vector Winds
6. Wrapping Up the Course

Where are we?

Bayesian Models

- Bayesian models differ from frequentist models only in that the parameters $\boldsymbol{\theta}$ are random.
- This allows us to stack priors to create hierarchical models.
- The Gibbs sampler is a universal algorithm that allows us to efficiently sample from the posterior in hierarchical models.

Example: Rater Model

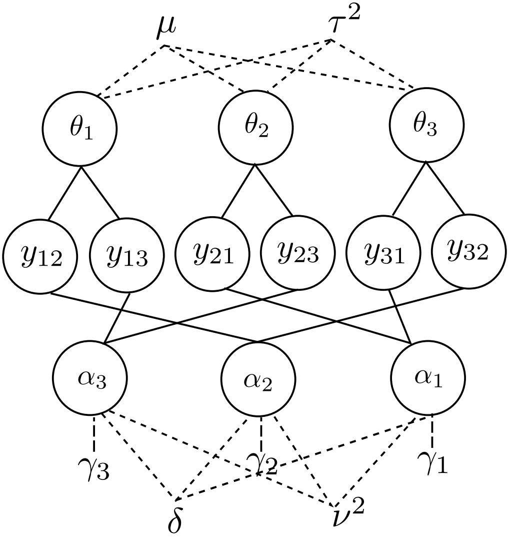

Computations

- The Gibbs sampler provides samples from the posterior.
- We can use these samples to estimate the posterior distribution (e.g., histogram) or the posterior mean.
- JAGS will automatically simulate from the posterior.

Results for the Rater Model

I put priors on γ_k and δ_ℓ so that I could do inference on them.

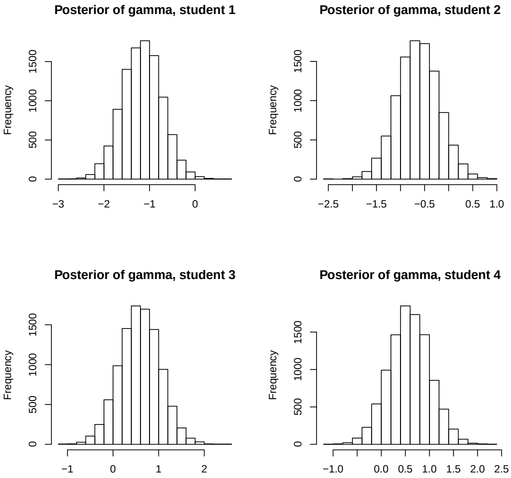

Where are we?

What would a frequentist do?

- **Remember:** In the Bayesian framework, to perform inference on a parameter, you must put a prior on it. Otherwise, you must specify its value beforehand.
- **Example:** Suppose there are 3 sections of a class taught by different professors. Let y_{ij} denote the final exam score of student j in class i.
  - $y_{ij} \big| \theta_i \sim N(\theta_i, \sigma^2)$ (instructor effect)
  - $\theta_i \sim N(\mu, \tau^2)$
- Maybe we can instead try to estimate *hyperparameters* such as μ and τ^2 from the data, then use $N(\hat{\mu}, \hat{\tau}^2)$ as the prior.
- What would our estimates of the parameters be then?

Empirical Bayes

- The idea of estimating hyperparameters from the data is called *empirical Bayes*.
- It is ultimately a frequentist method because we don't need to specify a prior on the hyperparameters we estimate!
- We get hierarchical models without subjective priors! Is this too good to be true?
- What are the challenges of doing this in practice?

Where are we?

Problem Setup

- El Niño is characterized by warmer sea surface temperatures (SST) in the equatorial Pacific Ocean.
- Therefore, to predict El Niño, one needs to forecast the SST several months in advance.
- We observe the average monthly SST at different locations in the Pacific:

$$\boldsymbol{z}_t \stackrel{def}{=} (\boldsymbol{z}_t(\boldsymbol{s}_1), ..., \boldsymbol{z}_t(\boldsymbol{s}_m))$$

- Want to forecast SSTs τ months in advance: $\boldsymbol{z}_{T+\tau}$.
- This analysis is taken from Berliner, Wikle, and Cressie (2000).

First Model: Linear Dynamical Model

$$\textbf{State model:} \quad \boldsymbol{y}_{t+\tau} = \Phi \boldsymbol{y}_t + \boldsymbol{\epsilon}_t \qquad \boldsymbol{\epsilon}_t \sim N(0, \Sigma)$$

$$\textbf{Data model:} \quad \boldsymbol{z}_t = A\boldsymbol{y}_t + \boldsymbol{\delta}_t \qquad \boldsymbol{\delta}_t \sim N(0, \sigma^2 I)$$

- A contains the first k principal components of the empirical covariance matrix over the spatial locations.
- $\boldsymbol{y}_t$ represents weights on those PCs.
- Unknown parameters are Φ, Σ, σ^2. Need to put priors on all of these:

$$\text{vec}(\Phi) \sim N(\text{vec}(0.9I), 100I)$$

$$\Sigma^{-1} \sim \text{Wishart}\left(\frac{1}{100(k-1)}I, k-1\right)$$

$$\sigma^2 \sim \text{InverseGaussian}(0.1, 100)$$

First Model: Linear Dynamical Model

1997 1998

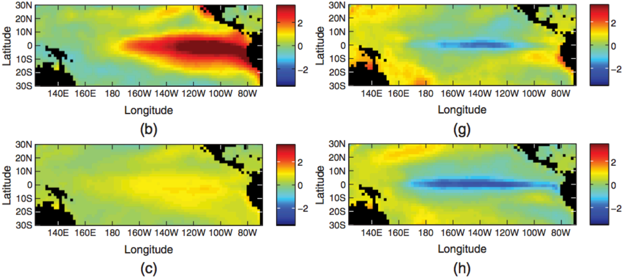

2.5% and 97.5% quantiles of posterior

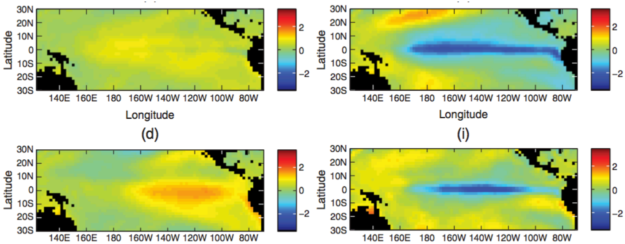

Second Model: "Non-Linear" Model

$$\textbf{State model: } \boldsymbol{y}_{t+\tau} = \Phi_t \boldsymbol{y}_t + \boldsymbol{\epsilon}_t \qquad \boldsymbol{\epsilon}_t \sim N(0, \Sigma)$$

$$\textbf{Data model: } \boldsymbol{z}_t = A\boldsymbol{y}_t + \boldsymbol{\delta}_t \qquad \boldsymbol{\delta}_t \sim N(0, \sigma^2 I)$$

- Allow $\Phi_t = \Phi(I_t, J_t)$ to vary with time.
- I_t classifies the current regime as "cool", "normal", or "warm". Obtained by thresholding the Southern Oscillation Index (SOI):

$$I_t = \begin{cases} 0 & \text{if } SOI_t < \text{low threshold} \\ 1 & \text{if } SOI_t \text{ in between} \\ 2 & \text{if } SOI_t > \text{upper threshold} \end{cases}$$

- J_t is obtained by similarly thresholding a latent process W_t:

$$W_t | \boldsymbol{\beta}, \tau^2 \sim N(\boldsymbol{x}_t^T \boldsymbol{\beta}, \tau^2)$$

Second Model: "Non-Linear" Model

1997 1998

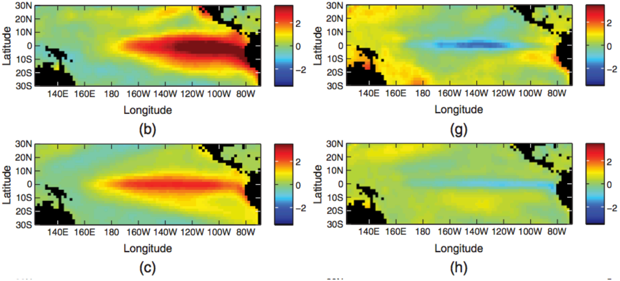

2.5% and 97.5% quantiles of posterior

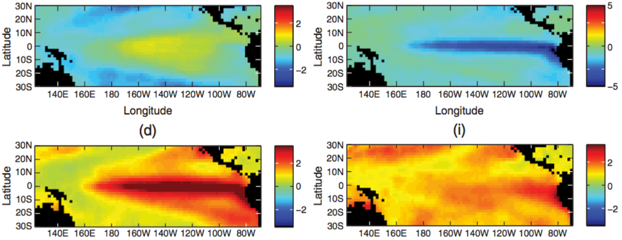

Where are we?

Problem Setup

- The Eurasian Collared Dove (ECD) was first observed in North America in the 80s and are now spreading quickly throughout the continent.
- They pose a threat to native ecosystems, so we would like to forecast their spread.
- Observe $z_t(\boldsymbol{s}_i)$: number of doves observed at location $\boldsymbol{s}_i$.
- This analysis is taken from Hooten, Wikle, Dorazio, and Royle (2007).

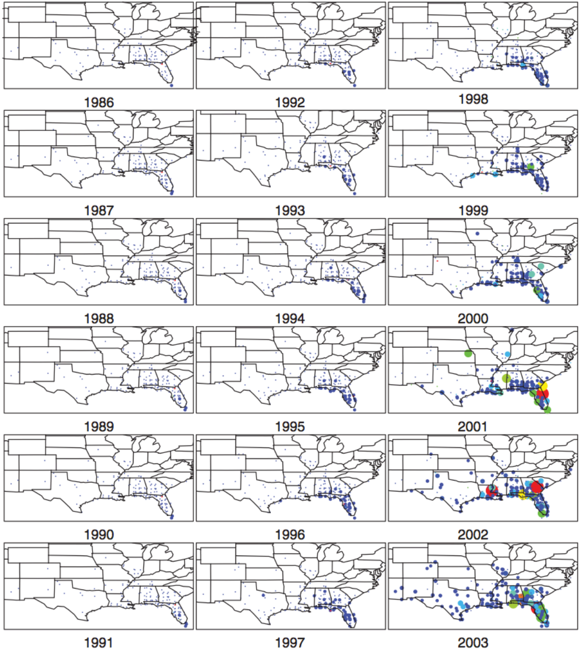
1986
1992
1998
1987
1993
1999
1988
1994
2000
1989
1995
2001
1990
1996
2002
1991
1997
2003

The Model

Data model: $z_t(\boldsymbol{s}_i) \mid y_t(\boldsymbol{s}_i), \pi \sim \text{Bin}(y_t(\boldsymbol{s}_i), \pi)$

$$\boldsymbol{y}_t \mid \boldsymbol{\lambda}_t \sim \text{Pois}(H\boldsymbol{\lambda}_t)$$

State model: $\boldsymbol{\lambda}_t = B(\boldsymbol{\alpha})G(\boldsymbol{\lambda}_{t-1}; \boldsymbol{\theta})\boldsymbol{\lambda}_{t-1}$

- π is the probability of observing an animal. Not estimable from this data alone, but the authors estimated it using data collected on a related species.
- G is a diagonal matrix that models growth, while B models dispersion.
- The authors go on to put priors on $\boldsymbol{\alpha}$ and $\boldsymbol{\theta}$.

Results

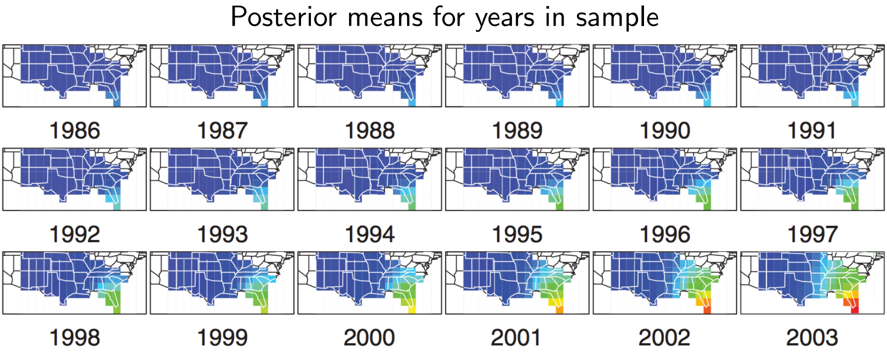
Posterior means for years in sample
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003

Results

Posterior means for future years

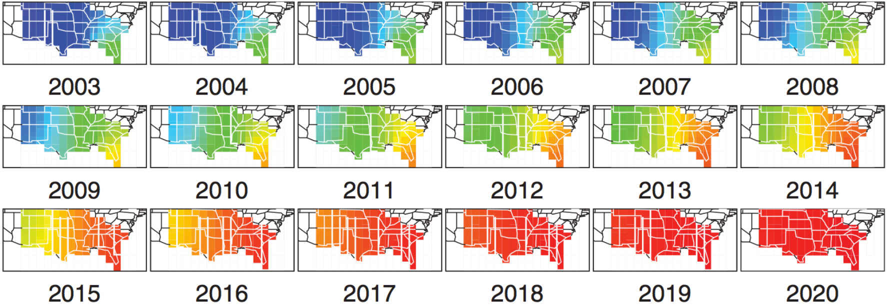

Where are we?

Problem Setup

- The goal is to predict wind speeds and directions at different locations over the Mediterranean.
- $(x_t(\boldsymbol{s}_i), y_t(\boldsymbol{s}_i))$ is the vector indicating the wind speed in the x- and y-directions.
- We assume that $\boldsymbol{x}_t$ and $\boldsymbol{y}_t$ are noisy measurements of underlying states $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$.

Data on February 1, 2005

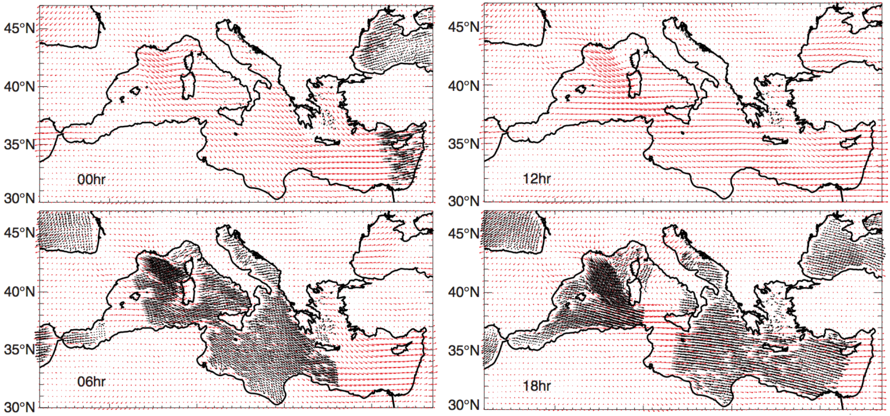

A Physical State Model

- The state model is movtivated by the Rayleigh friction equations:

$$\frac{\partial u}{\partial t} - fv = -\frac{1}{\rho_0}\frac{\partial p}{\partial x} - \gamma u \qquad \frac{\partial v}{\partial t} + fu = -\frac{1}{\rho_0}\frac{\partial p}{\partial y} - \gamma v.$$

  where u, v are the east-west and north-south components, p is the sea-level pressure, and the rest are (unknown) parameters.

- An approximate solution to these equations is given by

$$u \approx -\frac{f}{\rho_0(f^2+\gamma^2)}\frac{\partial p}{\partial y} - \frac{\gamma}{\rho_0(f^2+\gamma^2)}\frac{\partial p}{\partial x} - 2\frac{\gamma}{\rho_0(f^2+\gamma^2)}\frac{\partial u}{\partial t} - \frac{1}{\rho_0(f^2+\gamma^2)}\frac{\partial^2 p}{\partial x \partial t}$$

$$v \approx \frac{f}{\rho_0(f^2+\gamma^2)}\frac{\partial p}{\partial x} - \frac{\gamma}{\rho_0(f^2+\gamma^2)}\frac{\partial p}{\partial y} - 2\frac{\gamma}{\rho_0(f^2+\gamma^2)}\frac{\partial v}{\partial t} - \frac{1}{\rho_0(f^2+\gamma^2)}\frac{\partial^2 p}{\partial y \partial t}$$

- We can discretize this as follows:

$$\boldsymbol{u}_t = a_1 D_y \boldsymbol{p}_t + a_2 D_x \boldsymbol{p}_t + a_3 \boldsymbol{u}_{t-1} + a_4 D_x \boldsymbol{p}_{t-1}$$

$$\boldsymbol{v}_t = b_1 D_x \boldsymbol{p}_t + b_2 D_y \boldsymbol{p}_t + b_3 \boldsymbol{v}_{t-1} + b_4 D_y \boldsymbol{p}_{t-1}$$

Results for February 2, 2005

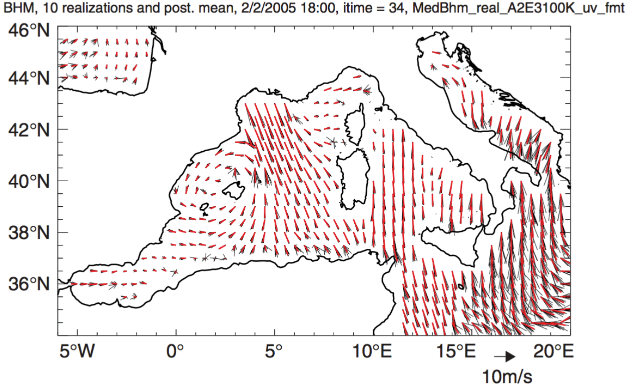

Where are we?

A Few Last Thoughts

- The analysis of spatial and temporal data is really the analysis of correlated data.
- "Model the mean function, use spatial and temporal methods to model the residual."
- There are two main ways to capture correlations: model the correlation directly (e.g., kriging) and via autoregressions.
- Many methods that work well for time series (e.g., Kalman filter) break down in space because the data are no longer ordered.

Thanks for a great quarter!