

# *A Markov random field spatio-temporal analysis of ocean temperature*

MICHAEL LAVINE\* and SUSAN LOZIER†

*Institute of Statistics and Decision Sciences, Duke University, 223 Old Chem Building, Durham, NC 27708*

Received April 1998; Revised December 1998

---

The National Oceanic Data Center (NODC) contains historical records from approximately 144,000 hydrographic stations in the North Atlantic. This data has been used by oceanographers to construct maps of point estimates of pressure, temperature, salinity and oxygen in the North Atlantic (Levitus (1994); Lozier *et al.* (1995)). Because data from any particular year are scarce, the previous maps have been for time-averaged values only. In addition, the maps have been reported without uncertainty estimates. This paper presents a Markov random field (MRF) analysis that can generate maps for specific time periods along with associated uncertainties.

To estimate changes in oceanic properties over time previous oceanographic work has focused on differences between a few time periods each having many observations. Due to data scarcity this poses a severe restriction for both spatial and temporal coverage of climatic change. The MRF analysis provides a means for temporal modeling that does not require high data density at each time period. To demonstrate the usefulness of a MRF analysis of oceanic data we investigate the temporal variability along 24.5°N in the North Atlantic. Our results are compared to an earlier analysis (Parrilla *et al.* (1994)) where data from only three time periods was used. We obtain a more thorough understanding of the temperature change found by this previous study.

*Keywords:* Bayesian analysis

1352-8505 © 1999  Kluwer Academic Publishers

---

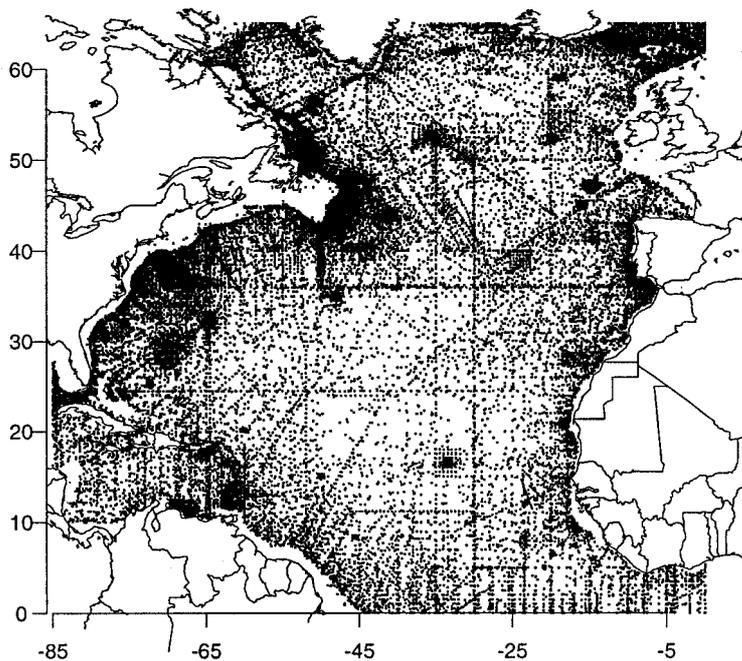
## 1. Introduction

The National Oceanic Data Center (NODC) collects data from oceanographic expeditions world-wide. This paper presents an analysis of ocean temperatures in the North Atlantic Ocean using data collected between 1905 and 1988, the most recent available to us at the time of writing. Fig. 1 shows the spatial locations of all stations (places where measurements have been made by oceanographic ships) in the North Atlantic. Fig. 2 shows the spatial locations and sampling dates of two data sets considered in this paper. Clearly there is no sampling regularity in either time or space.

Although the data are spatially distributed in three dimensions, a two-dimensional representation of the data is generally employed because properties in the ocean are, to first order, distributed by the mean flow and mixing processes along surfaces which are

\* Supported by NSF Grant DMS-9300137.

† Supported by NOAA Grant NA 46GP0184.

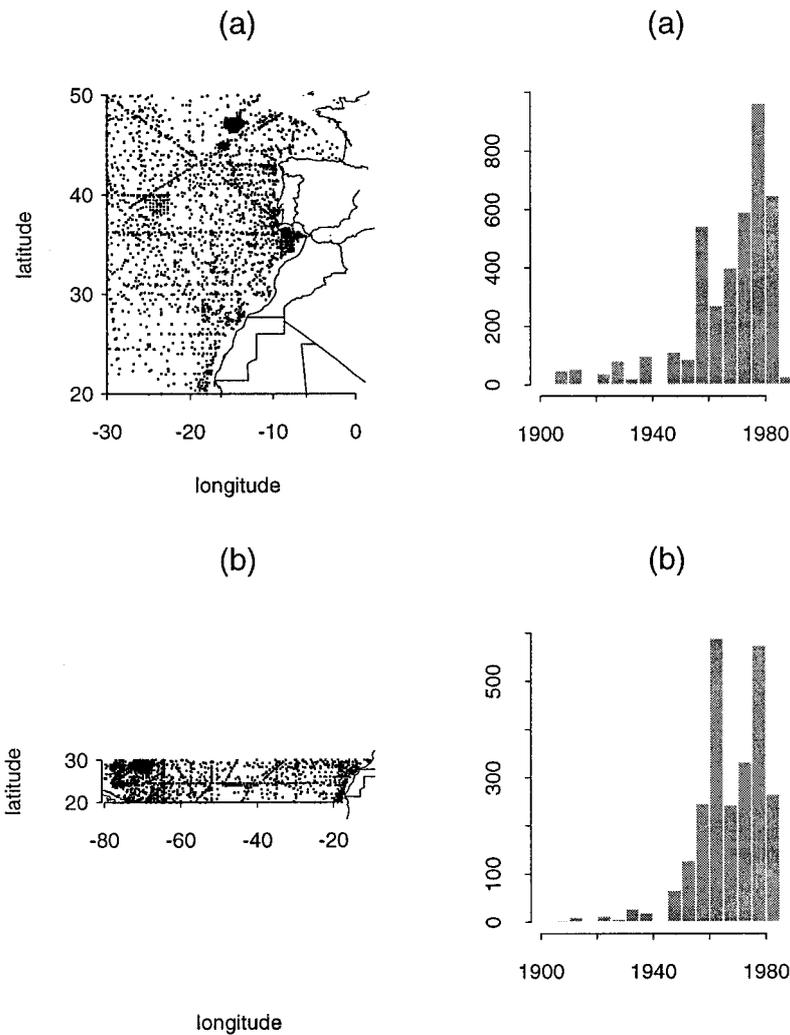


**Figure 1.** Data locations in the North Atlantic.

roughly horizontal. Specifically, oceanic flow and its associated properties tend to flow along surfaces of constant potential density (the density of the water after volumetric and heat effects due to compression have been removed), termed isopycnals (Lozier *et al.* (1994)). Fig. 3 shows a cross section of isopycnals in the North Atlantic at 24.5°N latitude. Isopycnals are not surfaces of constant depth; the isopycnals in Fig. 3 deviate from a horizontal plane, with shoaling at the eastern and western boundaries indicating boundary currents. In addition, isopycnals are not fixed in space. Each isopycnal may move up and down over time, reflecting changing oceanic conditions. Fig. 3 shows time-averaged locations of the isopycnals at 24.5°N latitude, averaged over the approximately 80 years represented in the NODC historical database.

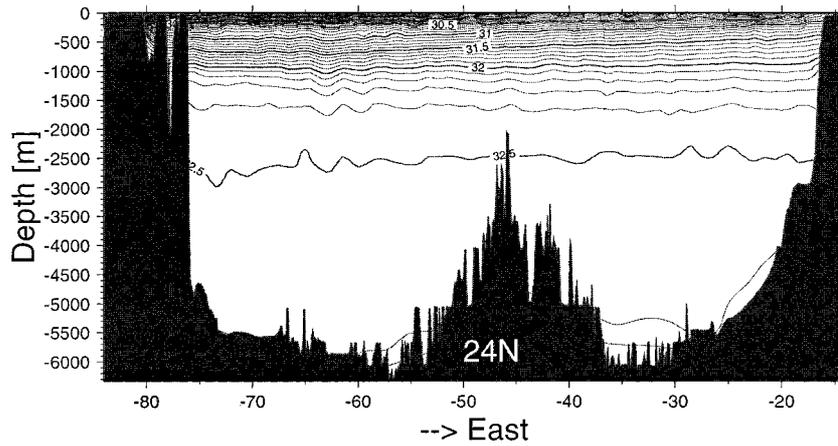
Our work presents an analysis of a single isopycnal, the  $\sigma_1 = 32.00 \text{ kg/m}^3$  surface ( $\sigma$  is calculated as  $\sigma = (\rho - 1) \times 1000$ , where  $\rho$  is the density of the water in cgs units.  $\sigma_1$  is the designation for potential density referenced to 1000 meters), which lies approximately at a depth of 1000 meters in the midlatitudes of the North Atlantic. The physical properties of ocean water at a given latitude and longitude are measured at discrete depths in the water column. When a temperature “measurement” is desired on an isopycnal that was not measured directly, the raw data are interpolated to that isopycnal. (Density is a function of pressure, temperature and salinity, which are all measured.) Interpolation error is expected to be small due to dense vertical sampling of the water column. For the purposes of this paper, we absorb the interpolation error into our model’s error term.

One important task for oceanographers is to estimate, along with a measure of accuracy, the temperature at locations in space-time for which there are no measurements. One



**Figure 2.** Locations and dates of two data sets. (a): near the Mediterranean; (b): 20°–30°N latitude.

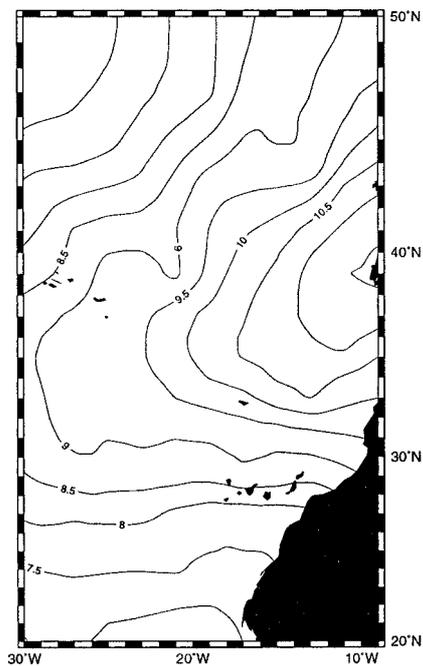
approach among oceanographers is exemplified by Levitus (1994) and Lozier *et al.* (1995). The raw data are first subject to quality control in which suspicious data points are deleted. The remaining data points are then smoothed locally to produce a reconstruction of the temperature surface which is often displayed as a contour map such as Fig. 4. The reconstruction takes no account of time and is treated as a “climatological mean” despite being heavily weighted to post 1955, as evident from the histograms in Fig. 2. From this, and similar reconstructions of pressure, salinity and other physical properties, oceanographers deduce the climatological flow field. A central improvement that our



**Figure 3.** Isopycnals in the Atlantic Ocean along 24.5°N latitude. The top of the figure is the ocean surface; the black mass at the bottom is the ocean floor; the vertical scale is depth in meters; the horizontal scale is degrees of longitude.

analysis will provide is a climatological mean for each time period separately and the ability to construct a climatological mean in which all years receive appropriate weight. See Section 4.1 for illustration.

Another important task for oceanographers is to characterize temporal changes in the



**Figure 4.** Ocean temperatures on the  $\sigma_1 = 32.00$  isopycnal estimated from post-1974 data, each point equally weighted.

ocean. Because of data scarcity temporal change is often assessed from repeated hydrographic sections, a recent example of which is the study by Parrilla *et al.* (1994). In this work transects across  $24.5^{\circ}\text{N}$  in the North Atlantic, taken in the years 1957, 1981 and 1992, were examined for changes in the temperature of the ocean waters. Importantly, their analysis used only the data from these three years and only along this one cross-section. Our work reanalyzes the temperature change along  $24.5^{\circ}\text{N}$ , incorporating information from neighboring latitudes and time periods. Our result confirms the change found by Parrilla *et al.* (1994). However, our more detailed modeling allows us to reconstruct the time series of temperature along  $24.5^{\circ}\text{N}$  and to see the observed temperature change in the context of the entire time series. Further, by examining isopycnals as well as isobars (surfaces of constant depth), we can begin to see some of the physical reasons for the temperature change. These issues are elaborated in Section 4.2.

Our overall goal is to provide estimates of observable ocean properties; this paper reports on temperature. In essence, we are using a Markov random field analysis to estimate the ocean's temperature on a regular grid in space and time from irregularly distributed measurements. Our premise is simple, namely, that our best estimate of an ocean temperature at a gridpoint derives from temperature measurements taken near that gridpoint in both space and time. Although our use of isopycnal surfaces involves a priori knowledge of the ocean state, we otherwise, purposely, use no physical dynamics in our model. Although it has become common to use observational data in conjunction with dynamical models to predict the ocean flow field (see, for example, Miller (1986), Fukumori & Malanotte-Rizzoli (1995) and Cane *et al.* (1996)), our goal is to provide an estimate of the temperature field, not the flow field.

A logical physical constraint on our temperature estimates could be a mathematical statement of the conservation of heat. However, the use of this conservation statement requires knowledge of the flow field and knowledge of the relative roles of advection (the transfer of a property by the large-scale flow field) and diffusion, both of which are imperfect at best. A later extension could be to incorporate conservation principles as a constraint on the estimated temperatures. But our analysis here is useful as it stands. The high data density of the North Atlantic makes our analysis more plausible in that there is a plentitude of data both in space and time in this heavily-sampled basin. In a basin where data coverage is sparse, such as the South Atlantic or South Pacific, the use of additional constraints may not be a choice, but a necessity. In sum, we are essentially creating an improved climatology which could, in turn, be used in data assimilation schemes to produce an improved estimate of the ocean flow field. In this paper, however, we confine ourselves to analyzing the temperature field itself.

## 2. The model

Our aim is to develop a model that accounts for spatial variability and that yields useful estimates of posterior uncertainty. We adopt a Markov random field (MRF) (Besag (1974)) model. The essence of MRF modeling is embodied in the following ideas.

(1) Every site  $i$  (in our application  $i$  is an ordered triple  $(j, k, l)$  indexing latitude, longitude and year), has a set of neighboring sites  $N(i)$ . If  $i' \in N(i)$  then  $i \in N(i')$  and we write  $i \sim i'$ .

(2) For every site  $i$ , let  $t_i$  be the value of the field (temperature, in our application) at that site. The conditional (prior or posterior) distribution of  $t_i$  given the  $t_{i'}$ 's at all other sites depends only on those in  $N(i)$ . That is,

$$p(t_i | t_{-i}) = p(t_i | \{t_{i'} : i \sim i'\})$$

where  $t_{-i}$  is the set of all temperatures except the  $i$ -th.

(3) The set of conditional distributions in item 2 completely determines (under fairly general conditions) the joint distribution of  $t$ .

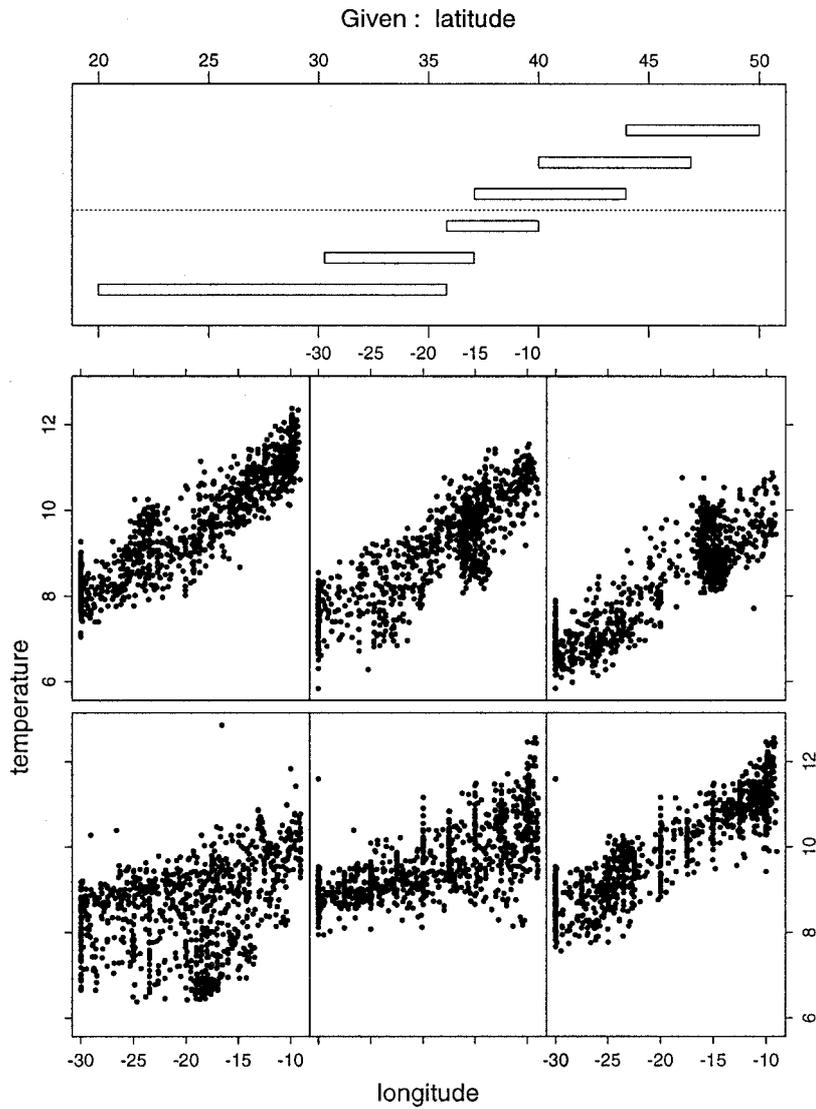
A more complete description can be found in Besag (1974), which first set out the MRF general theory and the conditions mentioned in item 3 above. A discussion of modeling and computational issues can be found in Besag *et al.* (1995), which illuminates the close connection between Markov Chain Monte Carlo (MCMC) computational methods and MRF spatial modeling and provides an excellent review of the two subjects. There is a large body of work using Markov random fields for spatial modeling. We mention here two recent papers, Waller *et al.* (1997) and Sun *et al.* (1997), which are concerned with spatial modeling of disease rates.

An apparent disadvantage of Markov random field modeling for data that arises at arbitrary locations in space and time is that the MRF requires dividing the domain into bins. An alternative that avoids binning is kriging, or one of its generalizations, which requires estimating or modeling a covariance function. See, for example, Handcock and Wallis (1994). But that disadvantage is outweighed, for us, by the advantage of the Markov property—each  $t_i$  is modeled as a function only of its neighbors. But before explaining the MRF prior in more detail, we examine exploratory plots of the data.

Figs. 5 and 6 are coplots of temperatures on the  $\sigma_1 = 32.00$  surface in the Mediterranean region (Fig. 2a) as a function of longitude for a fixed range of latitude (Fig. 5) or as a function of latitude for a fixed range of longitude (Fig. 6). The main feature in the region is a ridge of high temperature running roughly west-southwest from the Portuguese coast. This ridge is the so-called “Mediterranean tongue” caused by warmer Mediterranean waters flowing into and mixing with the colder open Atlantic waters. Fig. 4 shows the Med tongue on a contour plot.

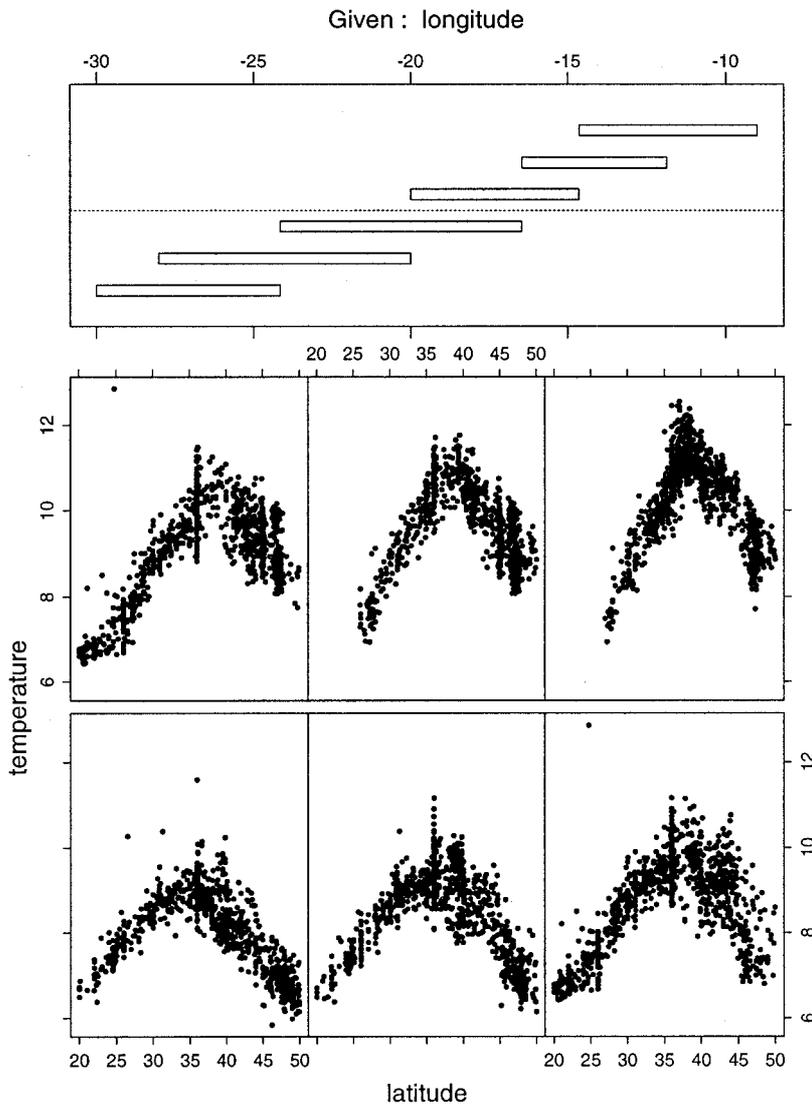
The temperature field in this region can be well approximated as locally linear, except possibly at the top of the ridge. Even there, if the bins are small enough, a locally linear approximation will likely work well. A local regression model (**loess** in **Splus**) estimates the residual standard error as 0.4044 in a locally quadratic fit and 0.4104 in a locally linear fit, confirming that the quadratic fit offers only modest improvement at best. Our MRF model corresponds roughly to a locally linear fit. In addition, none of the inferences presented in this paper depend crucially on the exact shape or location of the ridge.

Similar plots of temperatures on the  $\sigma_1 = 32.00$  surface in the  $24.5^\circ\text{N}$  region (Fig. 2b) suggest even more strongly that a locally linear model is adequate.



**Figure 5.** Coplot of temperature by longitude given latitude in the Mediterranean region. The *given panel* is at the top; the *dependence panels* are those in the  $2 \times 3$  array. On each dependence panel, temperature is graphed against longitude for those observations whose values of latitude lie in one of the intervals shown on the given panel. Moving from left to right through the given panel corresponds to moving from left to right and then bottom to top through the dependence panels (Chambers and Hastie (1992)).

Temperature measurements are taken at irregularly spaced locations in three dimensions: latitude, longitude and time. To implement the Markov random field, we impose a rectangular grid defined by  $J$ ,  $K$  and  $L$  evenly-spaced bins of latitude, longitude and time, respectively.



**Figure 6.** Coplot of temperature by latitude given longitude in the Mediterranean region. The *given panel* is at the top; the *dependence panels* are those in the  $2 \times 3$  array. On each dependence panel, temperature is graphed against latitude for those observations whose values of longitude lie in one of the intervals shown on the given panel. Moving from left to right through the given panel corresponds to moving from left to right and then bottom to top through the dependence panels (Chambers and Hastie (1992)).

Let  $N_{j,k,l}$  denote the number of measurements in bin  $(j, k, l)$ ,  $y_{j,k,l,n}$  (for  $n \in \{1, 2, \dots, N_{j,k,l}\}$ ) be the  $n$ -th such measurement and  $N \equiv \sum_{j,k,l} N_{j,k,l}$ . Let  $t_{j,k,l}$  denote the temperature in bin  $(j, k, l)$ , averaged with respect to time and space. In symbols,

$$t_{j,k,l} = \int_{\text{bin } j,k,l} t_{x,y,z}^{\text{true}} dx dy dz$$

where  $t_{x,y,z}^{\text{true}}$  is the true temperature at location  $(x, y, z)$ . Given  $t \equiv \{t_{j,k,l}\}$  and a data standard deviation  $\sigma^D$ , we model the data as

$$y_{j,k,l,n} \sim N(t_{j,k,l}, (\sigma^D)^2)$$

with the  $y_{j,k,l,n}$ 's being conditionally independent given  $(t, \sigma^D)$  and where  $\sigma^D$  accounts for measurement error, within-bin variation and vertical interpolation error. Based on our analysis of the historical data it is apparent that measurement and interpolation errors are ignorable compared to within-bin variation. We take a flat prior for the data precision  $\tau^D \equiv (\sigma^D)^{-2}$ , independent of the other parameters, because, with our large sample sizes the posterior from the flat prior will approximate well the posterior from any reasonable proper prior.

Our belief that temperature varies smoothly from place to place is modeled through a Gaussian pairwise difference Markov random field prior on  $t$ , which can be thought of as locally linear, or “the least squares fit of a plane to the values at the . . . neighbors” (Besag and Kooperberg (1995, p. 744)) and having “appeal if  $[t]$  is known to be a smooth surface” (Besag (1989, p. 398)), as we believe it is. For each pair of neighbors  $t_i$  and  $t_{i'}$  ( $i$  and  $i'$  are triples in which one coordinate has been incremented or decremented by 1.) we introduce a “precision”  $\tau_{i,i'} = \tau_{i',i} > 0$  and include in the prior the term

$$e^{-\frac{1}{2}\tau_{i,i'}(t_i - t_{i'})^2}.$$

The prior conditional mean of  $t_i$  given  $t_{-i}$  and the  $\tau_{i,i'}$ 's is a weighted average of the  $t_{i'}$ 's in  $N(i)$ , with weights given by the  $\tau_{i,i'}$ 's.

The usual formulation is that the  $\tau_{i,i'}$ 's are either a “set of specified nonzero weights” (Besag *et al.* (1995, p. 11)) or else all the same; i.e.,  $\tau_{i,i'} = \tau$  for all  $(i, i')$ , as in, for example, Waller *et al.* (1997) through their parameter  $\lambda$  (p. 610). But we are not prepared to specify the weights; we want to give them a prior. And the “lower-order homogeneous schemes, . . . , have been specifically designed with *local* stochastic interaction in mind; . . . it is unreasonable to apply them in situations where there is evidence of gross heterogeneity over the lattice” (Besag (1974, p. 205)). In our case, we don't have *a priori* evidence of gross heterogeneity, but we do want to allow for its existence by permitting the  $\tau_{i,i'}$ 's to vary spatially. We accomplish that by including the  $\tau_{i,i'}$ 's in the Markov random field.

The  $\tau_{i,i'}$ 's are of three types according to whether  $i$  and  $i'$  differ in the latitude, longitude or time coordinate. Within each type, we model spatial variation of the  $\tau_{i,i'}$ 's by adding to the prior the terms

$$e^{-\nu_c |\tau_{i,i'} - \tau_{j,j'}|}$$

for every neighboring pair  $(\tau_{i,i'}, \tau_{j,j'})$  where  $\nu_c$  is either  $\nu_{\text{lat}}$ ,  $\nu_{\text{lon}}$  or  $\nu_{\text{time}}$ .

Let  $\boldsymbol{\tau}$  be the collection of  $\tau_{i,i'}$ 's. Conditional on  $\boldsymbol{\nu} \equiv (\nu_{\text{lat}}, \nu_{\text{lon}}, \nu_{\text{time}})$ , the prior density is

$$\begin{aligned}
p(\tau^D, \mathbf{t}, \boldsymbol{\tau} | \mathbf{v}) \propto & k(\mathbf{v}) e^{-\frac{1}{2} \sum \tau_{i,i'} (t_i - t_{i'})^2} \\
& \cdot e^{-\nu_{\text{lat}} \sum |\tau_{i,i'} - \tau_{j,j'}|} \\
& \cdot e^{-\nu_{\text{lon}} \sum |\tau_{i,i'} - \tau_{j,j'}|} \\
& \cdot e^{-\nu_{\text{time}} \sum |\tau_{i,i'} - \tau_{j,j'}|}
\end{aligned}$$

where  $k(\mathbf{v})$  is a constant that depends on  $\mathbf{v}$  and where the first summation is over all neighboring pairs of temperatures and the second, third and fourth summations are over all neighboring pairs of “precisions” in the latitude, longitude and time directions, respectively. The use of absolute values of the pairwise differences between  $\tau$ 's is for computational convenience, as explained on page 19.

**Remark** This prior is improper on  $\tau^D$ , and also on  $\mathbf{t}$  because it only addresses pairwise differences between the  $t_i$ 's and not their overall level. See Besag (1989) for further comment. For fixed  $t_i$ 's (with  $t_i - t_{i'} \neq 0$ ), the prior is proper on  $\boldsymbol{\tau}$  because the first term acts as independent exponential priors for the  $\tau_{i,i'}$ 's.

In the prior, the overall level of  $\boldsymbol{\tau}$  is not well determined. But in the posterior the  $t_i$ 's will be known approximately and will serve to pin down the overall level of the  $\tau_{i,i'}$ 's and their pairwise differences.

**Remark** Allowing the  $\tau$ 's to vary arbitrarily in  $(0, \infty)$  is, as far as we know, a methodological innovation.

**Remark** If bins are of equal length in the latitude and longitude directions then one may wish to enforce  $\nu_{\text{lat}} = \nu_{\text{lon}}$ .

We have described a joint MRF prior for  $\mathbf{t}$  and  $\boldsymbol{\tau}$ . Each  $t$ , except for those on the boundary, has twelve neighbors: six other  $t$ 's and the  $\tau$ 's that join them. Each  $\tau$ , except for those on the boundary, has eight neighbors: the two  $t$ 's that it joins and the six nearest neighbor  $\tau$ 's.

### 3. Posterior calculations

The posterior is proportional to

$$\begin{aligned}
p(\tau^D, \mathbf{t}, \boldsymbol{\tau} | \mathbf{y}, \mathbf{v}) \\
\propto k(\mathbf{v}) (\tau^D)^{N/2} e^{-\frac{1}{2} \tau^D \sum_{j,k,l,n} (y_{j,k,l,n} - t_{j,k,l})^2 - \frac{1}{2} \sum_{i \sim i'} \tau_{i,i'} (t_i - t_{i'})^2 - \sum_c \nu_c \sum_{i,i' \sim j,j'} |\tau_{i,i'} - \tau_{j,j'}|}
\end{aligned} \quad (2)$$

and can be sampled by Markov Chain Monte Carlo (MCMC) methods. In Equation (2) the second summation is over all neighboring pairs  $t_i \sim t_{i'}$ ; the third is over all neighboring pairs  $\tau_{i,i'} \sim \tau_{j,j'}$ . See Besag *et al.* (1995) for a recent review of the method. We chose the particular MCMC algorithm known as Gibbs sampling (Gelfand and Smith (1990)) which Besag *et al.* (1995) and Sun *et al.* (1997) agree is well suited to this problem. Gibbs

sampling requires random variate generation from the so called complete conditional distributions—the conditional distributions of each variable given all the others. For our model these are

$p(\tau^D | \mathbf{y}, \mathbf{v}, \mathbf{t}, \boldsymbol{\tau})$ : This density is proportional to  $(\tau^D)^{N/2} e^{-\frac{1}{2}\tau^D \sum (y_{j,k,l,n} - t_{j,k,l})^2}$  and is therefore the  $\Gamma(1 + N/2, \frac{1}{2} \sum (y_{j,k,l,n} - t_{j,k,l})^2)$  density.

$p(\mathbf{t} | \mathbf{y}, \mathbf{v}, \tau^D, \boldsymbol{\tau})$ : This density is proportional to

$$e^{-\frac{1}{2}\tau^D \sum (y_{j,k,l,n} - t_{j,k,l})^2 - \frac{1}{2} \sum \tau_{i,i'} (t_i - t_{i'})^2}$$

and is therefore multivariate normal. Instead of drawing directly from the  $J \times K \times L$  dimensional density we draw from the univariate conditional distribution of  $t_{j,k,l}$  given all the other temperatures.

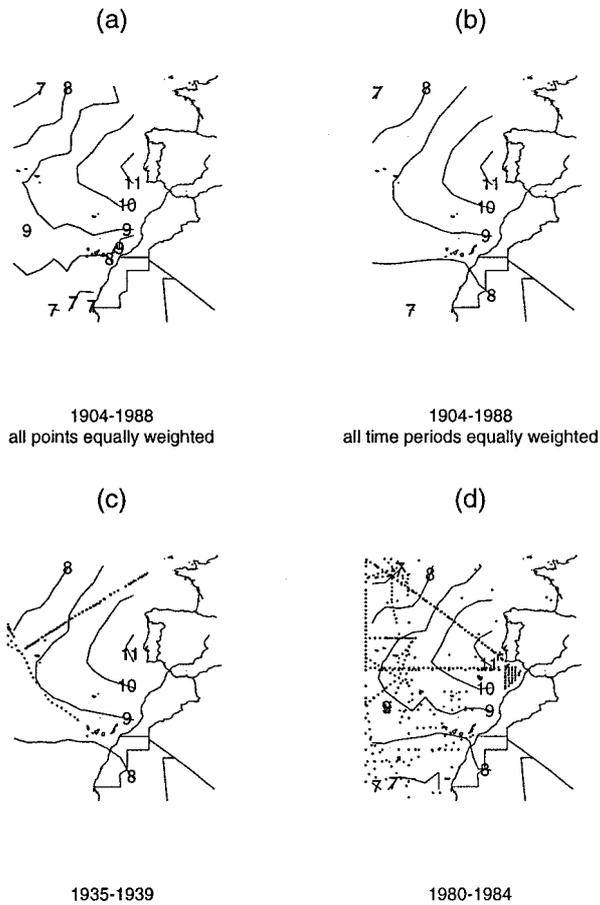
$p(\boldsymbol{\tau} | \mathbf{y}, \mathbf{v}, \tau^D, \mathbf{t})$ : Here again we draw from the distribution of  $\tau_{i,i'}$  given all the other  $\tau$ 's. The density is proportional to

$$e^{-\frac{1}{2}\tau_{i,i'}(t_i - t_{i'})^2 - \nu_c \sum |\tau_{i,i'} - \tau_{j,j'}|}$$

Let  $\tau_1 \leq \tau_2 \leq \dots \leq \tau_m$  be the ordered values of the  $\tau_{j,j'}$ 's that are neighbors of  $\tau_{i,i'}$ . Then the distribution is piecewise exponential on the intervals  $\{(0, \tau_1), (\tau_1, \tau_2), \dots, (\tau_m, \infty)\}$  and hence available for sampling.

The use of absolute values in the second term in the exponent is for computational convenience. It may appear that other choices are equally convenient but, for example, the use of squares would give a quadratic exponent  $\tau_{i,i'}(t_i - t_{i'})^2 - \nu_c \sum (\tau_{i,i'} - \tau_{j,j'})^2$  and hence a conditional Normal distribution with mean  $(2\nu_c \sum \tau_{j,j'} - (t_i - t_{i'})^2) / 2m\nu_c$  which may be arbitrarily far below 0. We judge such conditional distributions to be inappropriate for parameters restricted to  $(0, \infty)$  and so choose the absolute value alternative. While this choice does have implications for the smoothness of the  $\tau$  field (see Besag *et al.* (1995, pp. 11, 12) and references therein for discussion) we expect them to be unimportant as  $\tau$  is not a parameter of interest.

To assess convergence of the sampler we monitored four parameters of primary interest—average temperature in each of four time periods along the 24.5° transect from Africa to the Bahamas (details in Section 4.2)—and also, following Waller *et al.* (1997), “a representative subset (of size 24) of *all* the parameters”. In each case we used the values at every 500th iteration of the Gibbs sampler from the 10,000th through 60,000th. The run-time on our Digital AlphaStation 200 is about 3 to 4 days. CODA (Convergence Diagnosis and Output Analysis Software for Gibbs sampling output) (Best *et al.* (1995)) was used to calculate convergence diagnostics. Of the 4 monitored parameters of interest, all passed the Geweke test (with an initial bin fraction of 30%) and the Heidelberger and Welch test. Of the 24 other monitored parameters only one had a Geweke convergence diagnostic (Geweke (1992)) Z-score larger than 2 and one failed the Heidelberger and Welch (Heidelberger and Welch (1983)) test.



**Figure 7.** Reconstructions of the temperature field near the Mediterranean Sea. The dots in panels (c) and (d) are measurement locations.

## 4. Inferences

### 4.1 *The Mediterranean tongue*

Fig. 7(a) shows a reconstruction of the temperature field on an isopycnal in the Med region, done using the methodology described in Lozier *et al.* (1995). The reconstruction is called a “climatological mean” even though it weights all points equally and therefore overrepresents years with many data points. Unlike the Lozier *et al.* (1995) methodology, our analysis yields a posterior for each year separately. Figs. 7(c) and (d) show our posterior mean reconstructions for two of the  $L = 17$  time periods. With these posteriors in hand, one can construct climatological means in which time periods are weighted appropriately, perhaps equally or perhaps inversely by their posterior standard deviations. Fig. 7(b) shows one such reconstruction. It shows, for each (latitude, longitude) pair, the

average over time of the posterior mean temperature in each (latitude, longitude, time) bin, or, in symbols,

$$\sum_l \hat{t}_{j,k,l}/L$$

where  $\hat{t}_{j,k,l}$  is the posterior mean. In this case, there is little change through time, so Fig. 7(b) is similar in most respects to Fig. 7(a).

One difference between our analysis and that of Lozier *et al.* (1995) is that our reconstructions of different time periods have different posterior standard deviations. For example, the posterior standard deviations of the temperatures contoured in Fig. 7(b) is about  $0.2^\circ$ ; for Fig. 7(c) they range from about  $0.4^\circ$  to about  $0.6^\circ$ ; and for Fig. 7(d) they range from about  $0.2^\circ$  to about  $0.4^\circ$ . Another difference between Fig. 7(a) and our reconstructions is the slope of the eastern end of the  $8^\circ$  contour. The southward slope in our reconstructions is an edge effect from our model. To make serious inferences in that region we would modify the model.

## 4.2 From Africa to the Bahamas

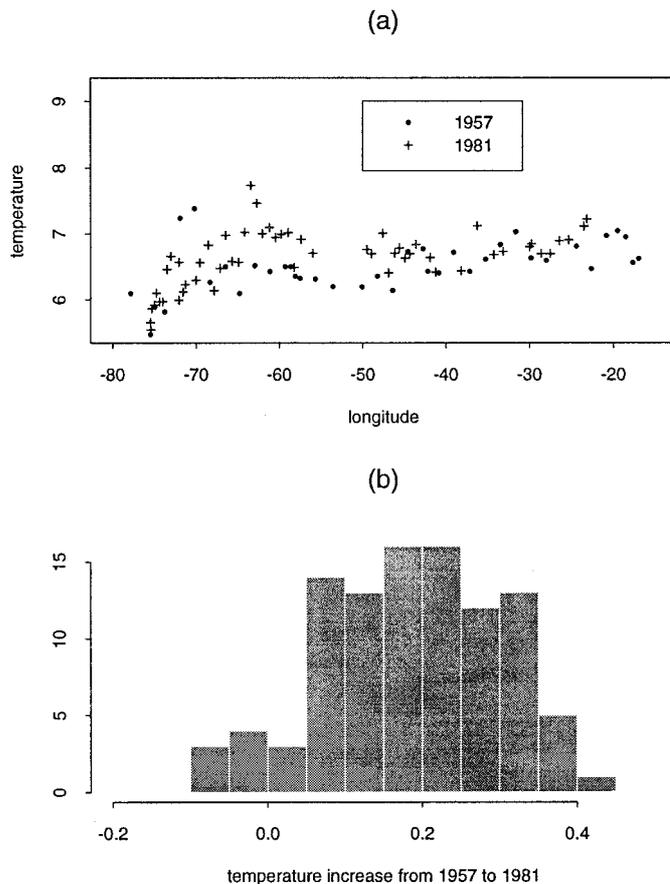
Parrilla *et al.* (1994) report on data taken along a transect of the Atlantic at latitude  $24.5^\circ\text{N}$  in 1957, 1981 and 1992. They

use the temperatures from all three surveys to show that the waters between 800 and 2,500 m depth have consistently warmed over the past 35 years and that the warming since 1957 is remarkably uniform across the east-west extent of the North Atlantic. The maximum warming, found at 1,100 m depth, is occurring at a rate of  $1^\circ\text{C}$  per century. The observed patterns of decadal-scale changes in ocean temperature are thus powerful signatures that can help us to understand the nature and causes of climate change.

Parrilla *et al.* (1994) find a warming of just over  $0.1^\circ\text{C}$  from 1957 to 1981, at around 1000 m depth, averaged across the transect. The 1957 and 1981 data are displayed in the first and last plots in the middle row of Fig. 12; the 1992 data have not yet been sent to the NODC and are not available to us. We use the Markov random field model to reanalyze the temperature change from 1957 to 1981 and address these questions:

- (1) Do we observe the same warming as Parrilla *et al.*?
- (2) Do results differ if the temperature is analyzed along an isopycnal rather than an isobar, as used by Parrilla *et al.*? (An isobar is a surface of constant pressure, or equivalently, constant depth).
- (3) How do the 1957 and 1981 data fit into a time series?

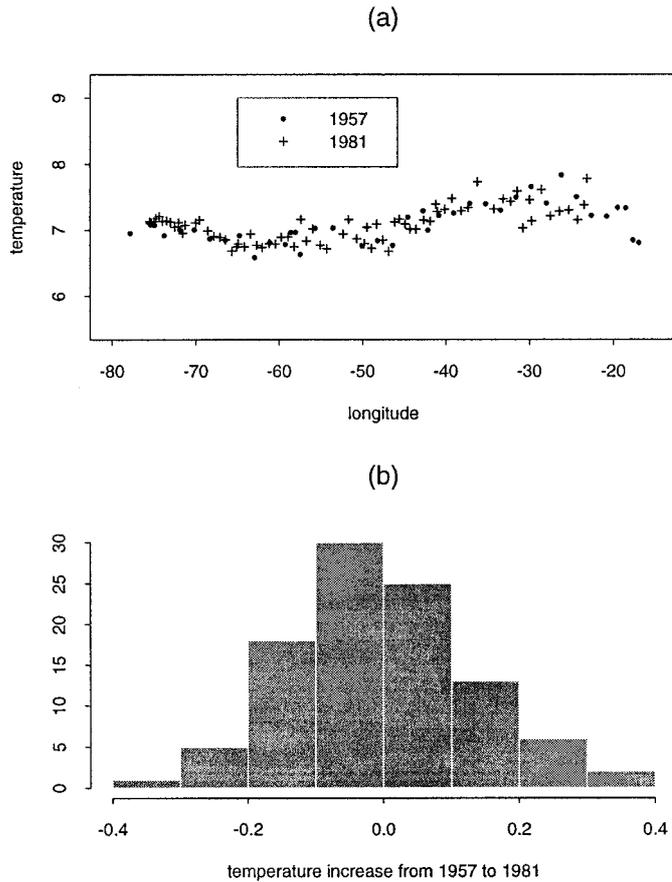
Because Parrilla *et al.* report warming on surfaces of constant depth, we answer the first question by running our model on data from the isobar at 1000 m depth. Each iteration of the Gibbs sampler gives a temperature realization in each bin. Within each iteration, averaging the bins at  $24.5^\circ\text{N}$  for 1957 and 1981 separately, and then subtracting, yields a



**Figure 8.** The  $24.5^{\circ}\text{N}$  transect. (a): the data; (b): posterior of temperature change from 1957 to 1981.

realization from our posterior distribution for the average temperature change. Fig. 8(a) shows the data collected in 1957 and 1981 between  $24^{\circ}\text{N}$  and  $25^{\circ}\text{N}$ . Except for two 1957 outliers, there is an apparent warming in the western basin along this surface. Fig. 8(b) shows a histogram of 100 realizations from the posterior for the average temperature increase along the transect. The mean temperature increase is 0.19 degrees; 93 of the realizations had temperature increases, 7 had decreases. Parrilla *et al.*'s reported  $0.1^{\circ}\text{C}$  warming is well supported by our posterior (as is any amount up to about  $0.35^{\circ}\text{C}$ ). Thus the answer to question 1 is that we do see the same warming as Parrilla *et al.*, at least qualitatively, although not precisely in all details.

As explained in Bindoff and McDougall (1994) and Bryden *et al.* (1996), there are two mechanisms to explain a warming (or any property change) at a particular depth in the ocean. Each mechanism has a different implication for the way in which the ocean is changing. Recall that the flow (and flow of heat) is along isopycnal surfaces. These surfaces are not fixed at a set depth, rather they may shoal or deepen depending on the local

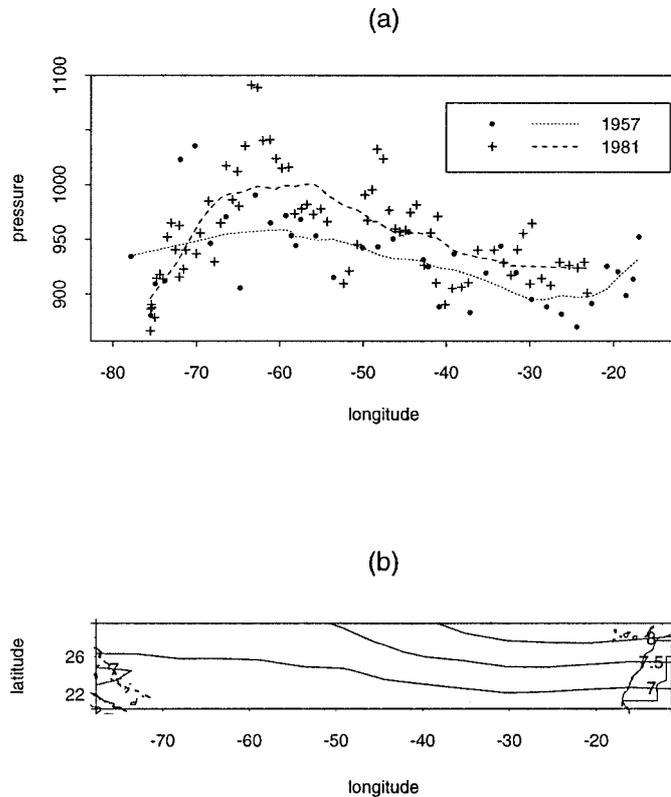


**Figure 9.** The 24.5°N transect. (a): the data; (b): posterior of temperature change from 1957 to 1981.

dynamics. If a warming at depth is observed one mechanism that can explain the change is that the depth of a particular isopycnal remained constant with time but the properties on the isopycnal changed. For the case at hand, the isopycnal in 1981 might be at the same depth as in 1957 but, at a given latitude and longitude, it might be warmer and saltier. The other mechanism is that the temperature of the isopycnal might remain constant but its depth may change over time. For the case at hand, the isopycnals in the intermediate depths of the North Atlantic may have become deeper than they were in 1957, thus bringing warmer waters to any fixed depth. Of course, these explanations are not mutually exclusive.

To distinguish between these possibilities we run our analysis on the isopycnal given by  $\sigma_1 = 32.00$ , which is approximately at 1000m depth near 24.5°N. (See Fig. 3.) This analysis, shown in Fig. 9, paints a different picture.

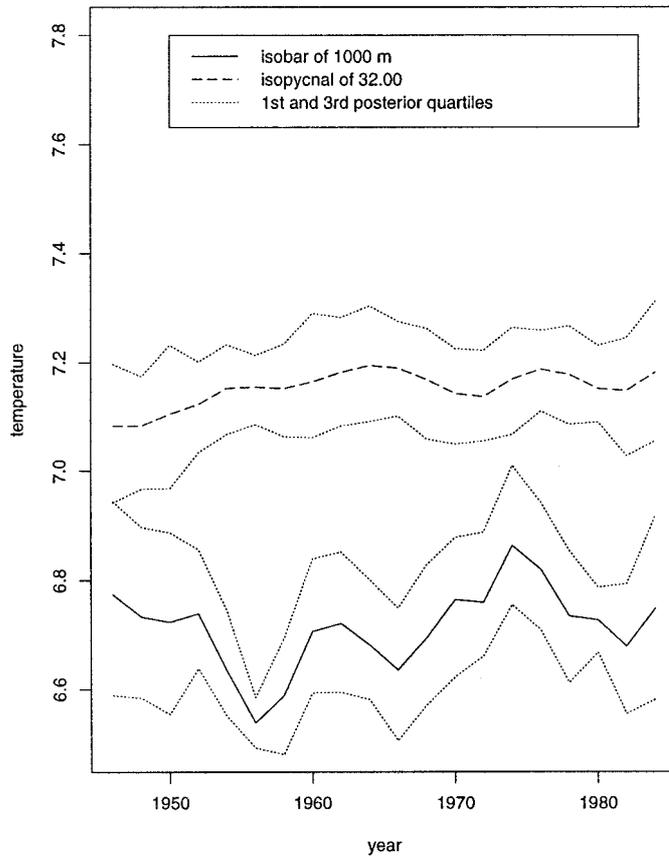
The posterior mean temperature increase, to two significant digits, is 0 and only 46 of the 100 realizations show any increase at all, suggesting that much of the warming along



**Figure 10.** The 24.5°N transect. (a): pressure on the isopycnal; (b): temperature contours.

the isobar is explained by a deepening of the isopycnal. Fig. 10(a), a plot of pressure (equivalently, depth) versus longitude, supports this view, showing that the 1981 isopycnal tends to be deeper than the 1957 isopycnal, as suggested by Joyce and Robbins (1996) and Bryden *et al.* (1996). This phenomenon is in accord with the supposition that the temperature contours (Fig. 10(b)) tend to meander north and south on time scales of less than a year. This horizontal meandering is manifest as a vertical meandering of the isopycnals past a fixed depth.

Finally, to answer question 3, we plot the posterior mean of the average transect temperature on both the isobar and the isopycnal in Fig. 11. We see that the average temperature of the isopycnal is relatively stable compared to the average temperature of the isobar, further supporting the idea that temperatures on the isopycnals have remained relatively constant but that the isopycnals meander vertically, causing the apparent temperature fluctuations on the isobar. The figure also indicates that Parrilla *et al.* were a bit unlucky in analyzing data from 1957, one of the coldest years since 1946. The result is that a more typical subsequent year, such as 1981, *appears* to be part of a warming trend. This illustrates, perhaps, the most important advantage of our analysis over traditional oceanographic methods: the ability to put the 1957 and 1981 data into a more complete



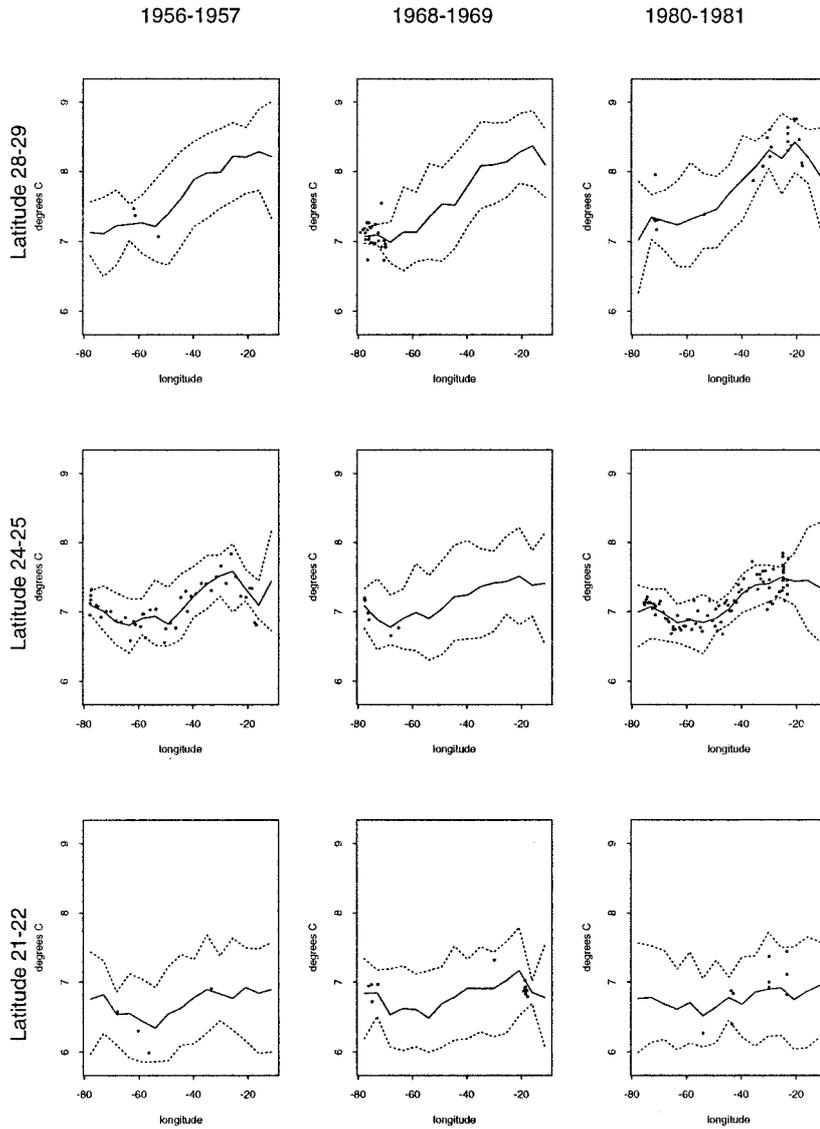
**Figure 11.** Time series plots of the average temperature along the 24.5°N transect for the isobar at depth 1000 m and the isopycnal  $\sigma = 32.00$ .

time series and to see that the 1957 to 1981 temperature increase may simply be due to natural decadal scale variability.

## 5. Model assessment

### 5.1 Model adequacy

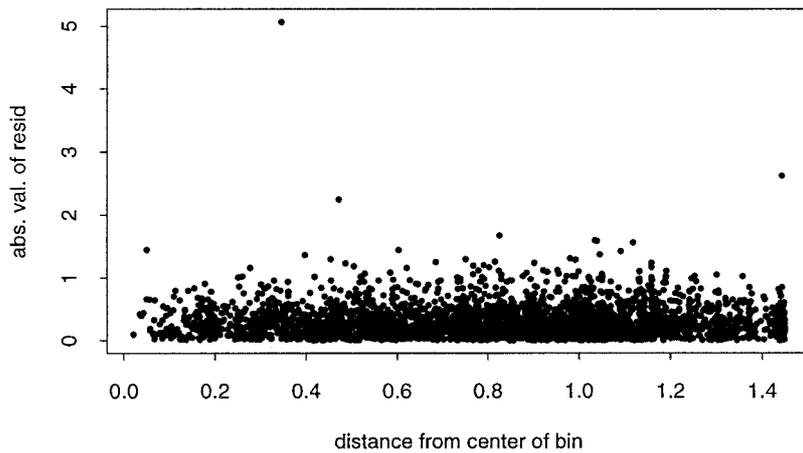
To check model adequacy we see how well it matches the data. Fig. 12 shows real data and posterior simulations for three latitudes and three time periods in the region of interest. The lines indicate the posterior mean and the central 80% of the posterior for the true temperature. Note in particular that the presence of within-bin variation means that these lines are not 80% prediction regions. To get prediction regions one must also account for the data standard deviation  $\sigma^D$ , whose posterior distribution is concentrated between 0.3 and 0.4°C.



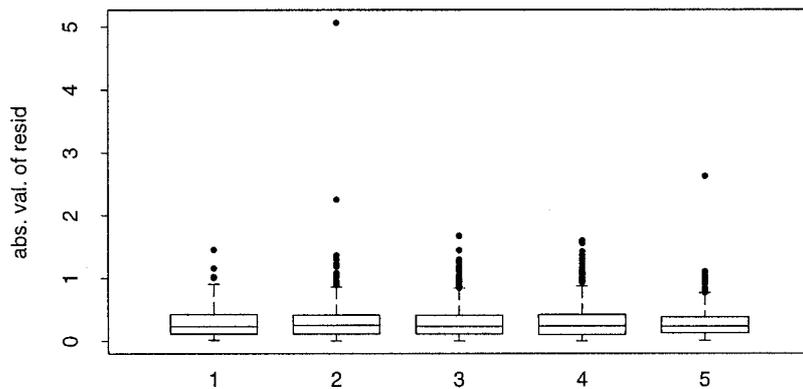
**Figure 12.** Simulated and measured temperatures as a function of longitude. Lines are 10th, 50th and 90th percentiles of simulations from the posterior.

Our model is perhaps overly simplistic in the way it uses local information. Three points in particular are worth noting.

- (1) We specify  $E[y_{j,k,l,n}|\mathbf{t}] = t_{j,k,l}$ . There might be some advantage to considering the exact location of the  $ijkl$ -th observation within the  $ijkl$ -th bin and modeling its mean as a function of  $t_{j,k,l}$  and the neighboring temperatures.



(a)



(b)

**Figure 13.** Residuals in the Med. tongue region as a function of distance from bin center. a) scatterplot. b) boxplots. The groups are quintiles based on distances from the centers of the bins.

(2) We use rectangular bins. There might be some advantage to hexagonal bins.

(3) We use a nearest neighbor scheme, corresponding to a locally linear fit. There might be some advantage to using more neighbors to yield a more flexible local fit.

For the Med region, where we used spatially square bins, Fig. 13 displays the absolute values of the residuals ( $y_{j,k,l,n}$  minus the posterior mean of  $t_{j,k,l}$ ) as a function of distance from the spatial center of the bin. The residuals are no larger nearer the edges of the bins

than near the centers. There is no evidence, therefore, that much is to be gained by more elaborate modeling of the mean.

## 5.2 Some $\nu$ considerations

Up to now we have treated  $\nu$  as a fixed parameter. A full Bayesian analysis would require solving for  $k(\nu)$ , which is problematic, so we chose instead initially to adopt a reasonable value for  $\nu$  and then examine sensitivity of inferences to that choice. The value of  $\nu$  determines the amount of smoothing in the three directions. To get an initial value we focussed on the data in Fig. 2(a).

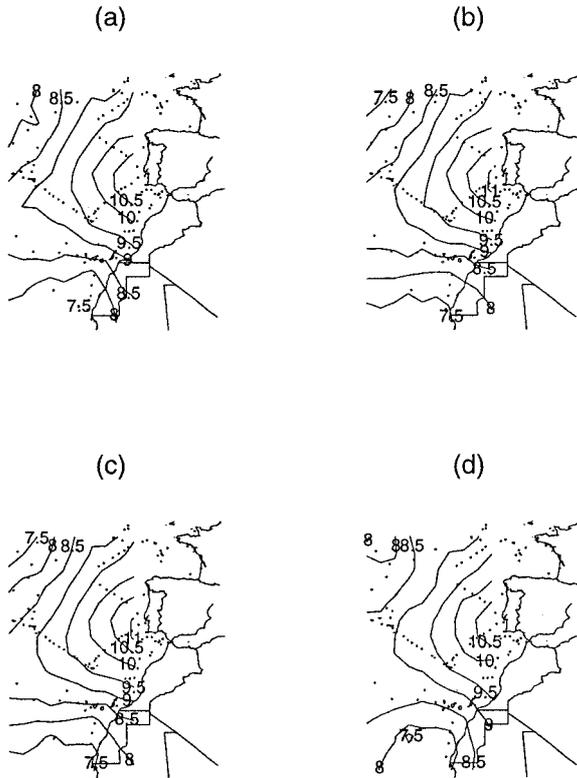
Fig. 4 shows a contour plot of temperature in the Mediterranean region calculated from the post-1974 data in the manner of Lozier *et al.* (1995). Because of the large number of points ( $N = 1638$ ) the contour plot gives a fairly accurate picture. We believe the Med tongue was also present in the early part of the century and therefore chose an initial value of  $\nu$  to yield what we believe to be a sensible posterior contour plot for the 1910–1914 temperatures—one in which the tongue is present, but slightly less well defined ( $N = 105$  in 1910–1914) than in Fig. 4.

Fig. 14 shows contour plots of posterior means for temperature in 1910–1914 each constructed from a different value of  $\nu$ . In each case, because we use spatially square bins, we take  $\nu_{\text{lat}} = \nu_{\text{lon}}$ . To begin, we note that the ridge of the Med tongue is very similar in all the reconstructions, running from a maximum of about 11°C to a minimum between 8.5°C and 9°C. The main differences between the reconstructions lie in the northwest and southwest corners where reconstructions (a) and (d) are warmer than (b) and (c). Reconstruction (a) is warm in those corners because  $\nu_{\text{time}}$  is large, which allows more flexibility in time, which in turn allows more spatial smoothing in the 1910–1914 time bin. Reconstruction (d) is warm due to the small values of  $\nu_{\text{lat}}$  and  $\nu_{\text{lon}}$ . Based on the oceanographic expertise of Lozier, we judge reconstructions (a) and (d) to be too warm in those corners, whereas reconstructions (b) and (c) appear to us to be sensible posteriors. So we choose an intermediate initial value of  $\nu = (0.5, 0.5, 0.00001)$ . We emphasize that we are *not* performing an empirical Bayes analysis and will *not* simply choose a single value of  $\nu$ . Rather, we make a careful analysis using an initial value of  $\nu$  and then check how sensitive our inferences are to that value. In Section 4.2 the initial value of  $\nu$  is adjusted to account for a change in bin size.

Since one of the key questions for oceanographers is the amount of warming on the isopycnal, we direct our sensitivity analysis to the histogram in Fig. 9. Fig. 15 shows posteriors, similar to Fig. 9(b), from four different analyses. Fig. 15(a) is, in fact, the same as Fig. 9(b). In each analysis we began with a 10,000 cycle burn in and then kept every 500th cycle of the next 50,000. Our conclusion is that the main result is insensitive to reasonable values of  $\nu$  and binsize: there is very little evidence of warming on the isopycnal and the warming on the isobar is due primarily to deepening of the isopycnal.

## 5.3 Other models

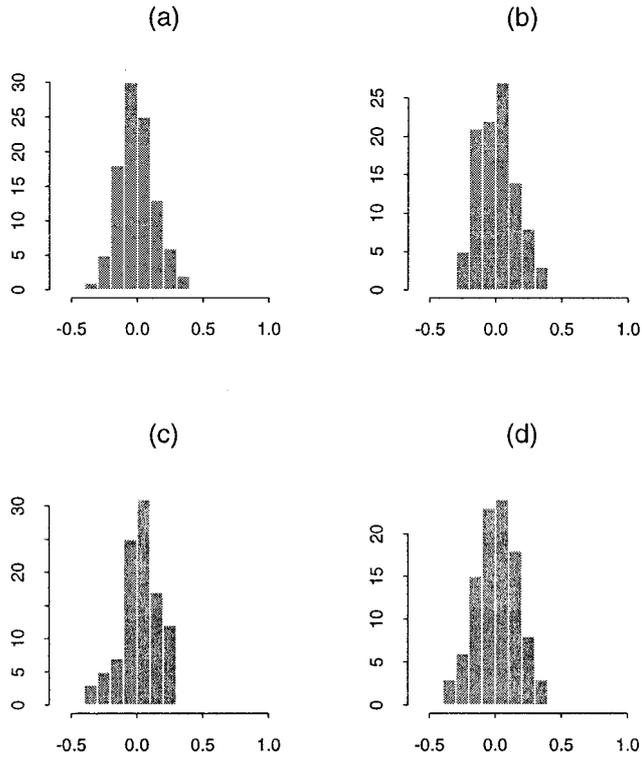
The value of our analysis depends on whether our numerical calculation of the posterior adequately captures the main features of the data set. It may fail to do so because either our



**Figure 14.** Contour plots of the posterior mean temperature in 1910–1914 using different values of  $\nu$ . The dots are locations of measurements taken between 1910 and 1914.

	$\nu_{\text{lat}}$	$\nu_{\text{lon}}$	$\nu_{\text{time}}$
(a)	0.5	0.5	0.02
(b)	0.5	0.5	0.0001
(c)	0.5	0.5	0.000001
(d)	0.05	0.05	0.000001

model or our calculations have led us astray. As a check on the main features of the posterior, we have also analyzed the  $24.5^\circ\text{N}$  isopycnal data using the local regression function `loess()` in **S-Plus**. We fit both a locally quadratic model and a locally linear model. For each model we computed the fitted temperature and standard error at latitude  $24.5^\circ\text{N}$ , at thirteen longitudes from  $15^\circ\text{W}$  to  $75^\circ\text{W}$ , in 1957 and 1981. We then subtracted the 1957 temperatures from the 1981 temperatures and averaged across the longitudes. We computed an approximate standard error by taking the square root of the sum of squares of the pointwise `se`'s. The average fitted temperature changes are  $0.12(\pm 0.14)$ , and  $0.07(\pm 0.08)$ , respectively (mean  $\pm$  `se`) and should be compared to the histogram in Fig. 9. We also used the two `loess` fits to predict the average transect temperature in 1946, 1948, ..., 1984. Those predictions are plotted in Fig. 16 along with the posterior mean from the



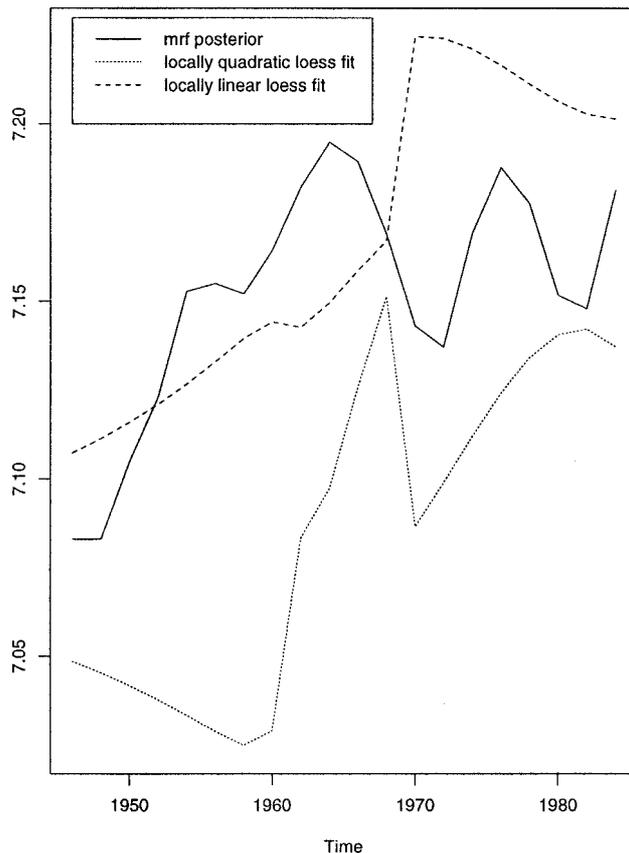
**Figure 15.** Four posteriors of isopycnal warming.

	$\nu_{\text{lat}}$	$\nu_{\text{lon}}$	$\nu_{\text{time}}$	$J$	$K$	$L$
(a):	0.2	0.7	0.000005	10	15	20
(b):	0.02	0.07	0.000005	10	15	20
(c):	0.2	0.7	0.0005	10	15	20
(d):	0.1	0.4	0.000003	20	30	40

MRF model. While not agreeing in detail, the analyses agree sufficiently well that we have no reason to mistrust either the MRF modeling or posterior calculations.

## 6. Summary

We have presented an analysis of ocean temperatures using a new Markov random field spatio-temporal model. We believe the model to be well suited to further analyses, either of temperature in other locations or of other oceanic properties such as salinity, pressure and oxygen concentration. The analysis makes three contributions to oceanography.



**Figure 16.** Three fits of temperature along  $24.5^{\circ}\text{N}$  at  $\sigma_1 = 32.00$ . The MRF posterior is the same as in Fig. 11.

- It gives reconstructions of the temperature field in time periods of sparse data.
- It facilitates a time series view of aggregate oceanic properties, such as the average temperature in a specified region.
- It yields what we believe are sensible estimates of the accuracy with which such reconstructions can be estimated.

Our use of the model for the  $24.5^{\circ}\text{N}$  transect revealed that the previously reported temperature increase from 1957 to 1981 may be part of natural decadal scale temperature variability and further, that the temperature variability along an isobar may be due to vertical movement coupled with temperature stability on isopycnals. Of course such change in the depths of a given isopycnal could also indicate changes in the climatic state of the ocean. For example, the depth of the isopycnals could be affected by long term changes in the surface wind forcing or by long term changes in the water mass structure in the tropical Atlantic. However, because the temperature changes on an isopycnal are relatively small, it is expected that the temporal variability of the isopycnal depth is approximately given by the temporal variability of the temperature on an isobar. The lower

curves in Fig. 11 suggest that the change from 1957 to 1981 can be as easily explained by natural decadal scale variability as by long term trend. This variability can be accounted for by the horizontal meandering of the thermal front in this locale.

One issue not yet discussed is the presence of ‘‘edge effects’’. That is, boundary sites are treated differently than interior sites by virtue of having fewer neighbors. This point is treated thoroughly by Besag and Kooperberg (1995). In our analysis of the Med region, we do not make inferences that involve edge sites critically. In the 24.5°N region, the average temperature along the transect includes sites on the eastern and western boundaries. To see whether edge effects might be coloring the inference, we recreated the histogram of Fig. 9(b), but excluding the edge sites. It made negligible visible difference.

Because our model is novel and complex, we have tried to be careful in implementation and interpretation. We checked convergence both numerically and graphically. We investigated sensitivity to the choice of the hyperparameter  $\nu$ . And we verified that a more standard local regression model yielded results similar to ours.

This raises the question, of course, why not simply use the standard model? For us the answers are first, that we believe posterior distributions are a more useful summary of uncertainty than are confidence intervals and second, we like the MRF approach. We intend, in future work, to continue investigating the ocean by implementing similar analyses on other ocean property fields, by modeling several isopycnals and several oceanic properties simultaneously and by investigating the possibility of incorporating conservation constraints into the model.

## References

- Besag, J., Green, P., Higdon, D., and Mengerson, K. (1995) Bayesian computation and stochastic systems. *Statistical Science*, **10**, 3–66.
- Besag, J. (1974) Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society (Series B)*, **2**, 192–225.
- Besag, J. (1989) Towards Bayesian image analysis. *Journal of Applied Statistics*, **16**, 395–407.
- Besag, J. and Kooperberg, Charles L. (1995) On conditional and intrinsic autoregressions. *Biometrika*, **82**, 733–746.
- Best, N.G., Cowles, M.K., and Vines, S.K. (1995) *CODA Manual version 0.30*, MRC Biostatistics Unit.
- Bindoff, N.L. and McDougall, T.J. (1994) Diagnosing climate change and ocean ventilation using hydrographic data. *Journal of Physical Oceanography*, **24**, 1137–1152.
- Bryden, Harry L., Griffiths, Michael J., Lavine, Alicia M., Millard, Robert C., Parrilla G., and Smethie, William M. (1996) Decadal changes in water mass characteristics at 24°N in the subtropical North Atlantic ocean. *Journal of Climate*, **9**, 3162–3186.
- Cane, Mark A., Kaplan, A., Miller, Robert N., Tang, B., Hackert, Eric C., and Busalacchi, Anthony J. (1996) Mapping tropical Pacific sea level: data assimilation via a reduced state space Kalman filter. *Journal of Geophysical Research*, **101**, 22599–22617.
- Chambers, John M. and Hastie, Trevor J. (1992) *Statistical Models in S*, Wadsworth & Brooks/Cole, Pacific Grove.
- Fukumori, I. and Malanotte-Rizzoli, P. (1995) An approximate Kalman filter for ocean data assimilation: an example with an idealized Gulf Stream model. *Journal of Geophysical Research*, **100**, 6777–6793.

- Gelfand, A.E. and Smith, A.F.M. (1990) Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association*, **85**, 398–409.
- Geweke, J. (1992) Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In *Bayesian Statistics (4)*, Valencia, Spain, Oxford University Press.
- Handcock, Mark S. and Wallis, James R. (1994) An approach to statistical spatial-temporal modeling of meteorological fields. *Journal of the American Statistical Association*, **89**, 368–378.
- Heidelberger, P. and Welch, P. (1983) Simulation run length control in the presence of an initial transient. *Operations Research*, **31**, 1109–1144.
- Joyce, Terrence M. and Robbins, P. (1996) The long-term hydrographic record at Bermuda. *Journal of Climate*, **9**, 3121–3131.
- Levitus, S. (1994) World ocean atlas 1994 CD-ROM sets. Informal Report 13, National Oceanographic Data Center.
- Lozier, M.S., McCartney, M.S., and Owens, W.B. (1994) Anomalous anomalies in averaged hydrographic data. *Journal of Physical Oceanography*, **24**, 2624–2638.
- Lozier, M.S., Owens, W.B., and Curry, R.G. (1995) The climatology of the North Atlantic. *Prog. in Oceanography*, **36**, 1–44.
- Miller, Robert N. (1986) Toward the application of the Kalman filter to regional open ocean modeling. *Journal of Physical Oceanography*, **16**, 72–86.
- Parrilla, G., Lavin, A., Bryden, H., García, M., and Millard, R. (1994) Rising temperatures in the subtropical North Atlantic Ocean over the past 35 years. *Nature*, **369**, 48–51.
- Sun, D., Tsutakawa, R.K., and Speckman, P.L. (1997) Bayesian inference for CAR(1) models with noninformative priors.
- Waller, Lance A., Carlin, Bradley P., Xia, H., and Gelfand, Alan E. (1997) Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association*, **92**, 607–617.

## Biographical sketches

Michael Lavine is Associate Professor of Statistics and Associate Professor of the Ecology Division of the Nicholas School of the Environment, Duke University. His current environmentally related work includes collaborations in water quality modeling, seedling mortality, hydrology, physical oceanography and the long-term effects of a CO<sub>2</sub> enriched atmosphere on forest growth.

Susan Lozier is Associate Professor of Earth and Ocean Sciences in the Nicholas School of the Environment, Duke University. Her research domain is physical oceanography, with a focus on the description and dynamics of large-scale ocean circulation, the modeling of cross-frontal exchange in coastal waters and the analysis of historical data to assess climatic trends in ocean properties.