Hierarchical Bayesian modeling of wind and sea surface temperature from the Portuguese coast

Ricardo T. Lemos*, Bruno Sansó † and F. D. Santos ‡

Abstract

In this work we revisit a recent analysis that pointed to an overall relaxation of the Portuguese coastal upwelling system, between 1941 and 2000, and apply more elaborate statistical techniques to assess that evidence. Our goal is to fit a model for environmental variables that accommodates seasonal cycles, long term trends, short term fluctuations with some degree of autocorrelation, and cross correlations between measuring sites and variables. Reference cell coding is used to investigate similarities in behavior among sites. We employ a Bayesian approach with a purposely developed Markov chain Monte Carlo method to explore the posterior distribution of the parameters. Our results substantiate most previous findings and provide new insight on the relationship between wind and sea surface temperature off the Portuguese coast.
Key words: upwelling, wind, SST, space-time models, Bayesian modeling

1 Introduction

The central task of climate change detection studies is to determine whether observed changes or trends in environmental time series are "significant", that is, highly unusual relative to the background of natural variability, and unlikely to have occurred by chance alone (Santer et al., 1996). Because most statistical models hinge on the assumption that the resulting residuals are independent and identically distributed (*viz.* Gaussian white noise), they are required to incorporate the most important sources of variability in the observed data; otherwise, this assumption is not verified and inference about the significance of trends is compromised. Commonly, environmental processes operate on various spatial and temporal

*Instituto de Oceanografia, Universidade de Lisboa, and Maretec - Instituto Superior Técnico, Universidade Técnica de Lisboa, Secção de Ambiente e Energia, Dpt. Mecânica, Av. Rovisco Pais, 1049-001 Lisboa, Portugal. E-mail rtl@net.sapo.pt

†Department of Applied Mathematics and Statistics, University of California, 1156 High St. MS:SOE2, Santa Cruz, CA-95064, U.S.A. E-mail bruno@ams.ucsc.edu, www.ams.ucsc.edu/∼bruno

‡SIM, Faculdade de Ciências da Universidade de Lisboa, Campo Grande, Edifcio C1, Piso 4, 1749-016 Lisboa, Portugal

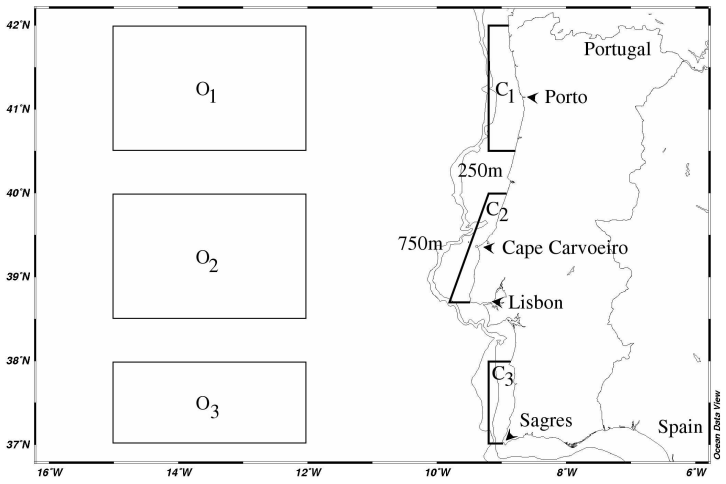

Figure 1: Location of weather stations and oceanic regions.

scales, producing time series that are impregnated with a number of features that make statistical modeling a challenging task. These include cycles, long-term linear and non-linear trends, short-term memory, spatial covariance and crossed covariance.

When starting from a simple statistical model that aims to detect long-term trends (*viz.* a model with intercepts, trends and uncorrelated Gaussian errors), red noise residuals are bound to result. In face of that, one of three approaches may be followed: i) modify the model, usually by making it more complex; ii) modify the input, by thinning the data set (Szunyogh et al., 2008) or pre-whitening (Rodionov, 2006); iii) modify the output, by correcting estimates based on the effective sample size (von Storch and Zwiers, 1999). At some point in i), parsimony becomes an issue, which may occur before the residuals conform to white noise. If the residual structure is small when compared to the modeled conterpart, "redness" is regarded as a nuisance property of residuals, and the modeler attempts to eliminate it by means of ii) or iii). In cases where residuals still contain relevant information, a different approach is needed to include it in the model. Such was, for example, the conclusion of Lemos and Pires (2004, henceforth, LP04), who used ordinary least squares regression to analyze wind and sea surface temperature data in the west Portuguese coast (37°-42° N, 9°-15° W) and found significant and readily interpretable amounts of information in the residuals.

This paper describes a more comprehensive method, based on hierarchical Bayesian modeling, which accommodates spatial, temporal and cross covariance structures, making statistical inference about trends and climatological cycles more accurate. The problem described by LP04 is revisited, so as to exemplify the method and compare results.

2 Model

Let

$$\overline{\boldsymbol{Y}}_t \equiv (\overline{V}_t(P), \overline{V}_t(CC), \overline{V}_t(L), \overline{V}_t(S), \\ \overline{SST}_t(O_1), \overline{SST}_t(O_2), \overline{SST}_t(O_3), \overline{SST}_t(C_1), \overline{SST}_t(C_2), \overline{SST}_t(C_3))'$$

denote a 10×1 vector containing monthly means of meridional wind speed (V) and sea surface temperature (SST), respectively collected at 4 weather stations of the Portuguese Institute of Meteorology – Porto (P), Cape Carvoeiro (CC), Lisbon (L) and Sagres (S) – and 6 oceanic regions off Portugal – $O_1, O_2, O_3, C_1, C_2, C_3$ (Figure 1). Regional SST means originated from raw data provided by the International Comprehensive Ocean-Atmosphere Data Set (Worley et al., 2005). Time t is written in months, from Jan-1941 (t=1) to Dec-2000 (t=720). Then,

$$\overline{\boldsymbol{Y}}_t \sim N_{10}\left(\boldsymbol{\mu}_t, \boldsymbol{\Omega}_t\right).$$

In this expression, the 10×10 covariance matrix $\boldsymbol{\Omega}_t$ is a diagonal matrix that accounts for uncorrelated measurement error; its i-th diagonal element is provided by $\omega(i)/n_t(i)$, where $\omega(i)$ and $n_t(i)$ respectively denote measurement error and the number of observations used to construct the monthly mean $\overline{Y}_t(i)$. The parameters $\omega(i)$, $i = 1, \ldots, 10$, are unknown; we set their prior distributions as inverse Gamma with mean 0.25 and scale parameter 1.

Regarding the 10×1 vector $\boldsymbol{\mu}_t$, we construct it as a sum of three components with distinct signals: one that accounts for the seasonal cycles and trends, $\boldsymbol{\theta}_t$; another that encompasses transient fluctuations thereof, $\boldsymbol{\lambda}_t$; and a third one composed of white noise with diagonal variance matrix $\boldsymbol{\Xi}$. Hence,

$$\boldsymbol{\mu}_t \sim N_{10}\left(\boldsymbol{\theta}_t + \boldsymbol{\lambda}_t, \boldsymbol{\Xi}\right).$$

The i-th diagonal component of $\boldsymbol{\Xi}$, $\xi(i)$, is an unknown parameter whose prior distribution is inverse Gamma with mean 0.025 and scale parameter 100. Following LP04, we describe $\boldsymbol{\theta}_t$ with form-free seasonal factors and linear trends,

$$\boldsymbol{\theta}_t \equiv \begin{pmatrix} \mathbb{1}_4 \alpha_m^{(V)} \\ \mathbb{1}_6 \alpha_m^{(SST)} \end{pmatrix} + \boldsymbol{\delta}_m + \left[\begin{pmatrix} \mathbb{1}_4 \beta_m^{(V)} \\ \mathbb{1}_6 \beta_m^{(SST)} \end{pmatrix} + \boldsymbol{\gamma}_m \right] (t - 360)/12,$$

where m indicates the month corresponding to time t. We also use reference cell coding (Kleinbaum et al., 1998), so that in the above expression, Sagres and the coastal region C_2 represent the "reference cells". This implies that the reference parameters, $\delta_m(S)$, $\delta_m(C_1)$, $\gamma_m(S)$ and $\gamma_m(C_1)$, are set to zero, for $m = 1, \ldots, 12$. Thus, $\alpha_m^{(V)}$ is the monthly intercept for wind, common to all weather stations; $\delta_m(P)$ is the difference in intercept between Porto and Sagres; $\beta_m^{(V)}$ is the monthly trend for wind, common to all weather stations; $\gamma_m(P)$ is the difference in trend between Porto and Sagres; and so forth for SST. Note that $\mathbb{1}_k$ denotes a $k \times 1$ vector of ones.

Finally, we describe $\boldsymbol{\lambda}_t$ by means of an autoregressive process of order one:

$$\boldsymbol{\lambda}_t = \mathrm{diag}(\boldsymbol{\rho})\boldsymbol{\lambda}_{t-1} + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim N_{10}\left(\mathbf{0}, \boldsymbol{\Sigma}\right), \tag{1}$$

The covariance matrix $\boldsymbol{\Sigma}$ follows, a priori, an inverse Wishart distribution with 26 degrees of freedom and scale matrix equal to 4 times the identity matrix. Diffuse Normal priors,

with mean zero and variance equal to 10^6, are used for most non-reference parameters. An exception is the 10×1 parameter vector $\boldsymbol{\rho}$, whose components have independent Uniform prior distributions with support $(-1, 1)$. Hence, the value of the components of $\boldsymbol{\rho}$ is, a priori, uniformly distributed inside the interval that corresponds to the assumption of stationarity of the autoregressive process.

To explore the posterior distributions of the parameters in the model described above, we employ Markov chain Monte Carlo (MCMC) methods. Gibbs sampling is applied to parameters with conjugate priors: the full conditional distributions (FCDs) of $\boldsymbol{\mu}_m$, $\alpha_m^{(V)}$, $\alpha_m^{(SST)}$ $\boldsymbol{\delta}_m$, $\beta_m^{(V)}$, $\beta_m^{(SST)}$, and $\boldsymbol{\gamma}_m$, $m = 1, \ldots, 12$, are Normals; the FCDs of $\xi(i)$, and $\omega(i)$, $i = 1, \ldots, 10$ are inverse Gammas; the FCD of $\boldsymbol{\Sigma}$ is inverse Wishart. The procedure to obtain these FCDs is standard; theoretical principles can be found *e.g.* in Gelman et al. (2003) and Gamerman and Lopes (2006), and a similar application to the one described above can be found in Lemos et al. (2007). To sample from each component of $\boldsymbol{\rho}$ at a time, we use a Metropolis-Hastings step, where the proposal distribution is a truncated Normal with $(-1, 1)$ support and variance set upon a pilot run. In the case of $\boldsymbol{\lambda}$, we apply the Forward Filtering, Backward Sampling method described by West and Harrison (1997, chapter 15) to the conditional multivariate dynamic linear model given by observation equation

$$\boldsymbol{\mu}_t - \boldsymbol{\theta}_t = \boldsymbol{\lambda}_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N_{10}\left(\mathbf{0}, \boldsymbol{I}\xi\right)$$

and evolution equation (1).

For convergence diagnostics, we employ the methods developed by Heidelberger and Welch (1983), Gelman and Rubin (1992), Geweke (1992), Raftery and Lewis (1992b), Raftery and Lewis (1992a), and Brooks and Gelman (1998), available in the package Bayesian Output Analysis Program (`BOA`) (Smith, 2005) within `R` (R Development Core Team, 2005). We use the default values of `BOA` to define the length of the burn-in stage, thin the chain, check stationarity and define the adequate sample size to achieve the precision required, when sampling from the posterior distribution.

In order to assess goodness of fit, we perform a suite of tests, including visual comparisons of time series of observations versus model estimates and associated 95% credibility intervals, informal inspections of red noise remaining in the model residuals ($\overline{\boldsymbol{Y}}_t - \boldsymbol{\mu}_t$), and a more systematic analysis of the upper level residuals, $\boldsymbol{\mu}_t - (\boldsymbol{\theta}_t + \boldsymbol{\lambda}_t)$, based on the ideas of Kim et al. (1998). Let Θ denote the collection of all model parameters; for each site s, the one step ahead distribution of $\mu_t(s)$ is

$$u_t^{[s]}(x) = P\left(\mu_t(s) \leq x | \Theta, x_i(s), i = 1, \ldots, t-1\right) \ .$$

Following Rosenblatt (1952), $u_1^{[s]}(\mu_1(s)), \ldots, u_{599}^{[s]}(\mu_{599}(s))$ are independent and uniformly distributed, provided the underlying distribution is continuous. Gneiting et al. (2005) denote this transformation as Probability Integral Transform and give an extensive list of references regarding its application.

We notice that, conditional on Θ, the one step ahead predictive distribution of $\mu_t(s)$ at any given site is normal. So, from each iteration of the MCMC after convergence, we can obtain a collection of random variables that should be independent and uniformly distributed. As in Kim et al. (1998), we consider a transformation to normality given by $\Phi^{-1}(u_t^{[s]})$, where

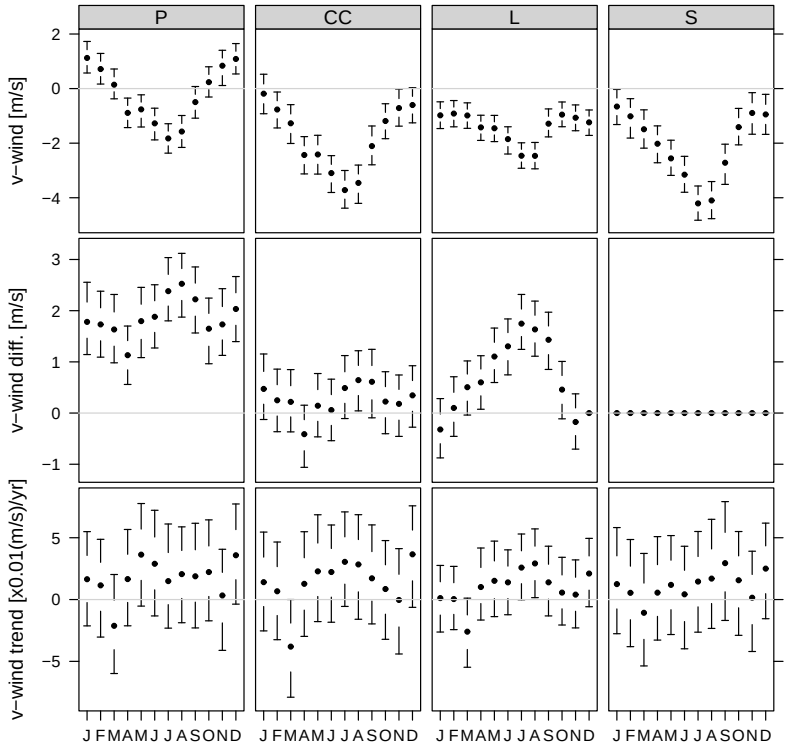

Figure 2: Posterior means (dots) and 95% credibility intervals (whiskers) for: the climatological v-wind cycle, i.e., $\alpha_m^{(V)} + \delta_m$ (upper row); the deviation in the cycle with respect to Sagres, i.e., δ_m (mid row); long term trends, i.e., $\beta_m^{(V)} + \gamma_m$ (lower row).

Φ^{-1} stands for the inverse of the standard normal cumulative distribution function. We use quantile-quantile plots, correlograms and periodograms to check that these variables are, respectively, normally distributed and independent.

3 Results

The seasonal cycle of the meridional wind component (v-wind), at the four weather stations analyzed, is depicted in the top panels of Figure 2. The cycle differs markedly from what a combination of a few harmonics would produce, making the form-free approach more adequate. Northerly, upwelling favorable winds predominate, especially between April and September; hence, we call this semester the climatological upwelling season. When compared to Sagres, Cape Carvoeiro displays nearly the same cycle, Porto shows a year-round shift towards more southerly winds, and Lisbon has a less pronounced seasonality (Figure 2, mid panels). We point out that the diurnal cycle of wind speed may be confounding this comparison, because the time of day at which wind measurements are taken differs among stations; see LP04 for a more detailed analysis.

The lower panels of Figure 2 present the long-term linear trends. Although the uncer-

Table 1: Posterior means and 95% credibility intervals for v-wind trends [$\times 10^{-2}$m s^{-1}yr^{-1}] and SST trends [$\times 10^{-2}$ °C yr^{-1}] during the upwelling season (April through September).

	P	CC	L	S	O_1	O_2	O_3	C_1	C_2	C_3
Upper Endpoint	4.25	4.33	3.24	3.72	1.11	1.18	1.41	2.53	2.73	2.46
Mean	2.27	2.23	1.80	1.38	0.11	0.30	0.44	0.94	1.34	1.15
Lower Endpoint	0.22	0.37	0.54	-1.42	-0.91	-0.68	-0.60	-0.78	-0.07	0.00

tainty associated with these estimates is considerable, some patterns can be discerned. The first is that, during the upwelling season, the trend appears more or less constant and positive, for all stations. On average, this rate is close to 0.0236 m s^{-1} yr^{-1} (Table 1), which corresponds to a weakening of 1.4 m s^{-1} in the 60 year period under study. The second noticeable feature is that, unlike other months, March presents a negative trend. Thus, toward the end of the century, northerly winds became on average stronger during March and weaker between April and September. December also presents a trend towards more southerly flow.

Figure 3 presents results for SST. In the reference region, C_2, the average annual SST is 16.4°C, and the difference between winter and summer months is close to 4.1°C. The seasonal cycle in the other two coastal regions is deviated by roughly 1°C, a finding that reflects the difference in latitude. Offshore regions also display the latitudinal effect, but their cycle is substantially different: if we compare C_2 with O_2, for example, we see that offshore waters can be warmer by more than 2°C , in late summer. This is clearly the effect of upwelling, which keeps coastal waters cooler than their offshore conterparts nearly year-round. Unlike wind, the seasonal cycles are fairly regular, so the form-free SST cycles seem amenable to be replaced with more parsimonious combinations of a few harmonics and an annual mean.

Similarly to v-wind, SST trend estimates display wide posterior 95% credibility intervals, especially along the coastline, where sampling is poorer. Nonetheless, coastal and offshore regions present distinct patterns: in the former, strong positive trends are frequent and negative trends are rare; in the latter, positive but weak trends emerge mostly in the first half of the year and negative trends are more common. Because posterior intervals for trends overlap substantially, we could perform the analysis with a single year-round trend, as LP04 did, or compute seasonal averages without great loss of information. Table 1 shows mean SST trends during the upwelling season. Even with a 6 months average, uncertainty regarding the magnitude of coastal warming is still considerable: taking C_2 as example, SST change from 1941 to 2000 ranges between -0.07°C and 1.64°C, with 95% confidence. The only region where the 95% posterior interval is within $\Re^+$ is C_3. In any case, we may state that overall, SST tended to increase along the Portuguese coast, rendering coastal-offshore gradients smaller. Simultaneously, the alongshore wind component weakened, from a starting point that was clearly upwelling favorable. ¿From these two observations, we may infer that the western Iberian upwelling regime relaxed towards the end of the 20$^{\text{th}}$ century. This conclusion is in agreement with previous findings, using both environmental (Lemos and Sansó, 2006; Álvarez et al., 2008) and biological (Lima et al., 2007a,b; Álvarez-Salgado et al., 2008) data.

It is worth noting that, along this coastline, typical summertime upwelling pulses last less

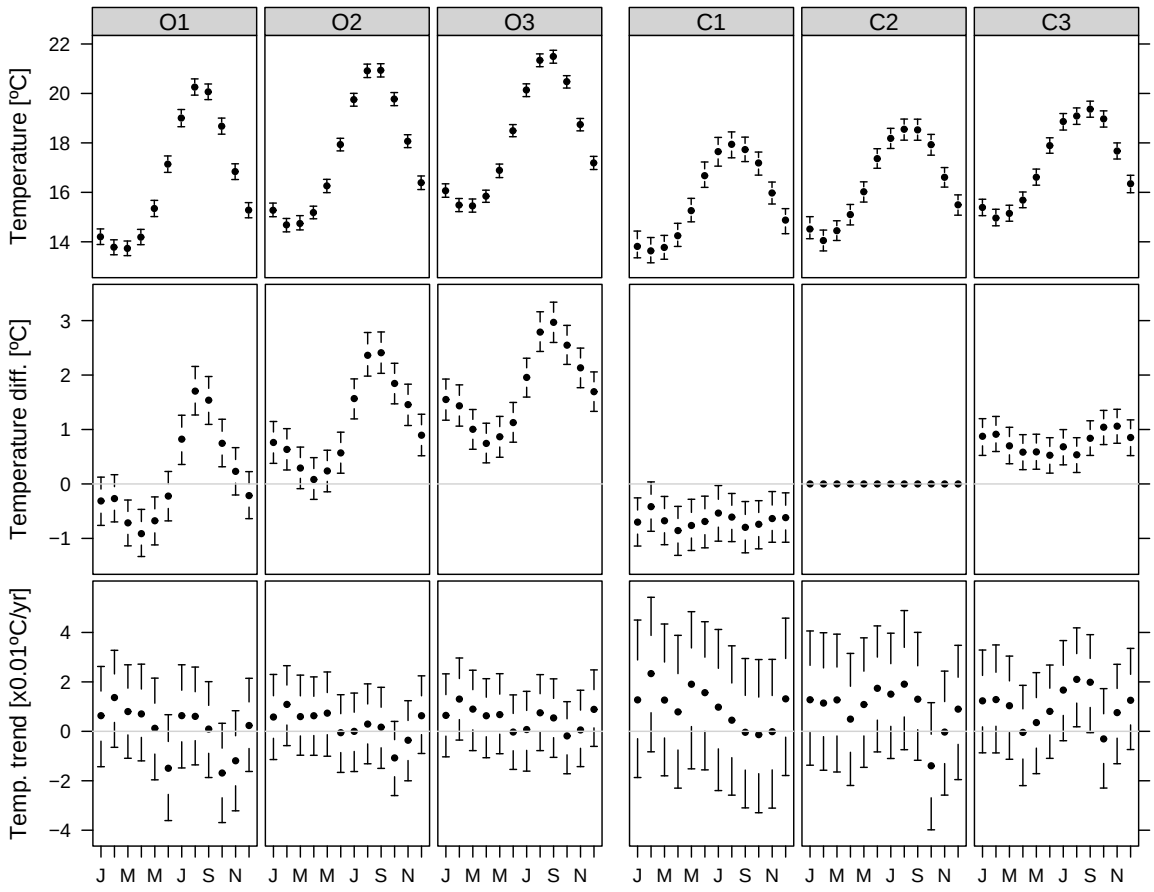

Figure 3: Posterior means (dots) and 95% credibility intervals (whiskers) for: the climatological temperature cycle, i.e., $\alpha_m^{(SST)} + \delta_m$ (upper row); the deviation in the cycle with respect to C_2, i.e., δ_m (mid row); long term trends, i.e., $\beta_m^{(SST)} + \gamma_m$ (lower row).

than one month (generally, about one week), and that two upwelling events may be parted by strong southerly wind conditions. Also, upwelling develops along a coastal strip that is sometimes narrower that the width of the coastal regions defined in this work. Hence, the present model operates on temporal and spatial scales that do not resolve individual upwelling events, but rather their net effect on monthly and regional means. As Figures 2 and 3 demonstrate, the signature of upwelling is still evident at these scales.

Finding similar climatologies and trends for the v-wind component (Figure 2) was a sensible result, given the spatial proximity of the weather stations under study and knowing that surface air flow along the Portuguese coast is mostly meridional. In line with this, it is also not surprising to find, as LP04 did, that high frequency deviations from the linear trends display some degree of spatial coherence. Namely, LP04 showed that these deviations – in their approach, residuals – are positively correlated among stations, and that nearby stations display stronger correlations than those farther apart. In the present work, we accommodate for this structure in the model, by using the matrix $\boldsymbol{\Sigma}$. Because we are using the same data set as LP04, we opt not to provide the above information in the prior distribution of $\boldsymbol{\Sigma}$. Rather, we assign few degrees of freedom to the prior and thus let the observations predominate in the posterior distribution of $\boldsymbol{\Sigma}$. To facilitate comparisons with LP04, we derive the correlation matrix from $\boldsymbol{\Sigma}$ and depict its posterior mean in a level plot. As Figure 4 shows, the correlation structure between weather stations mimics the features described by LP04.

The above discussion also applies to SST. In this part of the NE Atlantic, ffshore surface currents are mostly meridional, while upwelling/downwelling events induce similar degrees of cooling/warming along the coastline. Therefore, along-shelf SST covariances (i.e., offshore-offshore and coastal-coastal covariances) should be positive and stronger than their cross-shelf counterparts. Mesoscale ocean processes, on the other hand, should weaken covariances as distance increases. Figure 4 displays these features, thus confirming and extending the results of LP04.

One final noticeable feature presented in the correlation matrix, not analysed by LP04, is a block of positive correlations between wind and coastal SST fluctuations. This means that, when northerly winds (coded as negative v-wind values) are stronger than average, then coastal SST tends to drop below average, and vice versa. Because the association with offshore SST is weaker, we may interpret this as another sign of upwelling/downwelling on the analyzed time series.

As expected, transient v-wind fluctuations display weaker memory than SST. Sagres is the odd station, with strong lag-1 autocorrelation (Figure 5, left panel). As discussed in LP04, this may be a real feature or an artifact due to temporary instrument miscalibration. Regarding SST anomalies, the 95% posterior intervals concentrate probability between 0.4 and 0.5, and it is not possible to tell whether memory differs markedly between coastal and offshore regions. Finally, the central and right panels in Figure 5 depict the posterior distributions of the variance parameters ξ and ω. These distributions differ substantially among stations and regions, meaning that the model was able to learn from the data.

Residual analyses conformed well with the assumption that deviations between $\overline{\boldsymbol{Y}}_t$ and $\boldsymbol{\mu}_t$, as well as between $\boldsymbol{\mu}_t$ and $\boldsymbol{\theta}_t + \boldsymbol{\lambda}_t$, are consistent with white noise. We defer the interested reader to `www.ams.ucsc.edu/~bruno/windSST/`, where residual autocorrelograms,

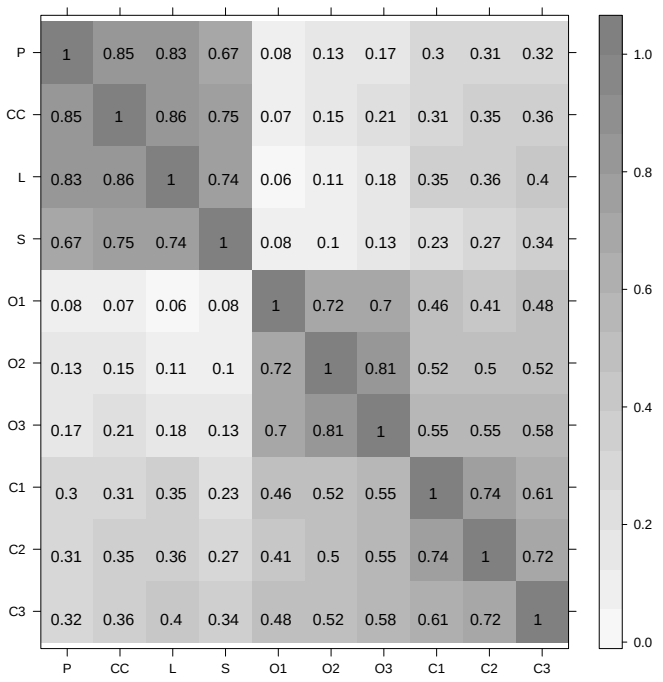

Figure 4: Posterior mean correlation matrix derived from the covariance matrix $\mathbf{\Sigma}$.

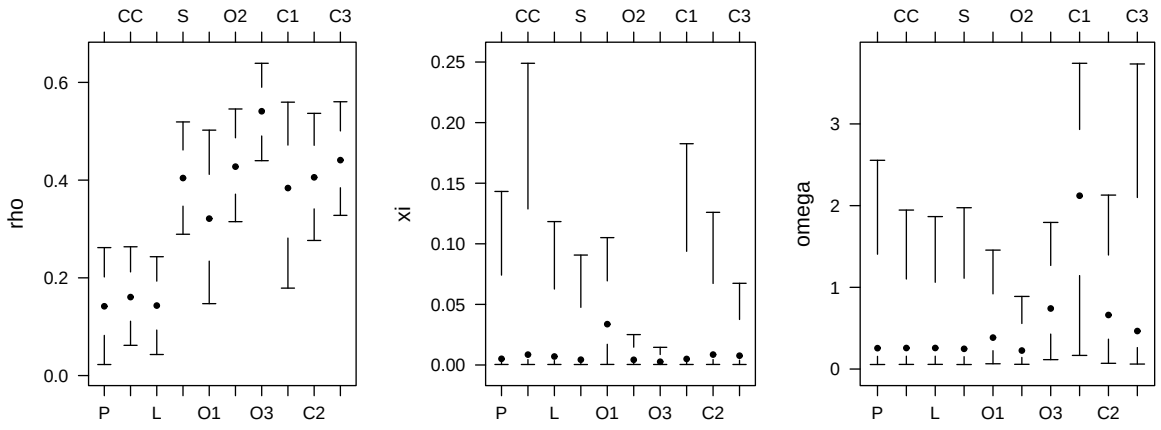

Figure 5: Posterior means (dots) and 95% credibility intervals (whiskers) for ρ (left panel), ξ (central panel) and ω (right panel).

cross-correlograms and normal quantile-quantile plots are available. MCMC convergence diagnostics can also be found there.

4 Final remarks

In this paper we present a statistical approach for the decomposition of environmental time-series into a number of parameters of interest, with the valuable asset that reliable measures of the estimation uncertainty are provided. The list of parameters consists of: a) stationary, form free seasonal cycles (α); b) long-term linear trends (β); c) terms for the comparison of cycles and trends among stations and regions (respectively, δ and γ); d) auto-correlation and cross-covariance parameters (respectively, ρ and Σ); e) error variances (ω and ξ). We use well established methods for the assessment of parameter convergence and goodness-of-fit analyses.

In an attempt to compare the approach with a more standard one, we revisit the problem described by LP04. Using their work as exploratory data analysis, we inflate the dimensionality of the model (item d in the list above), so as to provide more insight on the nature of the analyzed variables. Unlike LP04, we do not use a suite of hypothesis tests to reduce the dimensionality of the problem where possible; rather, we prefer to keep all model parameters and discuss their usefulness in capturing variability in the time series. With a Reversible Jump MCMC technique (see *e.g.* Hopcroft et al., 2007), we could have investigated the performance of nested models, at the cost of rendering the discussion more elaborate.

Overall, mean estimates of trends, seasonal cycles, auto-correlations and cross correlations are in good agreement between this paper and LP04. The correlation found between transient fluctuations of the v-wind component and coastal SST (around 0.30) provides further credit to the hypothesis that, even at the coarse spatial and temporal scales described, the link between the two variables is still strong enough to reflect upwelling. Hence, we substantiate the hypothesis forwarded by LP04, that a relaxation of upwelling off the west Iberian coast occured between 1941 and 2000.

On the other hand, the two papers differ markedly in how they treat uncertainty. Because parameter estimation is performed in a single modeling step, the present analysis allows an exchange of information between all types of parameters. In contrast, LP04 assumed that residuals from the "trend plus intercept" model consisted of white noise, so as to estimate confidence intervals for the trends and cycles, and later estimated residual auto-correlations and cross-covariances conditional on the mean intercepts and trends (i.e., assuming no uncertainty existed in these parameters). The procedure of LP04 thus violated two assumptions, with the resulting effect that confidence intervals were unduely narrow. The consequence for climate change detection is important: for instance, regarding SST trends, LP04 found evidence of significant offshore warming, while in the present paper and Lemos and Sansó (2006), it is not possible to ascertain whether this warming is spurious or not.

5 Acknowledgments

This work was performed under the scope of project Portcoast (POCTI /CLI /58348 /2004). R. Lemos acknowledges grant SFRH /BD /17929 /2004 from Fundação para a Ciência e a Tecnologia. B. Sansó was partially supported by the National Science Foundation grant NSF-Geomath 0417753.

References

Álvarez, I., Gomez-Gesteira, M., Decastro, M., and Dias, J. M. (2008). Spatiotemporal evolution of upwelling regime along the western coast of the Iberian Peninsula. *Journal of Geophysical Research - Oceans*, 113, C7, C07020, 10.1029/2008JC004744.

Álvarez-Salgado, X. A., Labarta, U., Fernández-Reiriz, M. J., Figueiras, F. G., Roson, G., Piedracoba, S., Filgueira, R., and Cabanas, J. M. (2008). Renewal time and the impact of harmful algal blooms on the extensive mussel raft culture of the Iberian coastal upwelling system (SW Europe). *Harmful Algae*, 6(6):849–855.

Brooks, S. and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7:434–55.

Gamerman, D. and Lopes, H. F. (2006). *Markov Chain Monte Carlo - Stochastic Simulation for Bayesian Inference*. Chapman and Hall, London, UK, second edition.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2003). *Bayesian Data Analysis, Second Edition*. Chapman & Hall/CRC.

Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7:457–72.

Geweke, J. (1992). Evaluating the accuracy of sampling-based i approaches to calculating posterior moments. In Bernardo, J. M., Berger, J. O., Dawid, P., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 7*. Clarendon Press, Oxford, UK.

Gneiting, T., Balabdaoui, F., and Raftery, A. (2005). Probabilistic forecasts, calibration and sharpness. Technical Report 483, Department of Statistics, University of Washington.

Heidelberger, P. and Welch, P. (1983). Simulation run length control in the presence of an initial transient. *Operations Research*, 31:1109–44.

Hopcroft, P. O., Gallagher, K., and Pain, C. C. (2007). Inference of past climate from borehole temperature data using Bayesian Reversible Jump Markov chain Monte Carlo. *Geophysical Journal International*, 171(3):1430–1439.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: likelihood inference and comparison with arch models. *Rev. Fin. Stud.*, 65:361–393.

Kleinbaum, D. G., Kupper, L. L., Muller, K. E., and Nizam, A. (1998). *Applied regression analysis and other multivariable methods*. Duxbury: Pacific Grove, CA.

Lemos, R. T. and Pires, H. O. (2004). The upwelling regime off the west portuguese coast, 1941-2000. *International Journal of Climatology*, 24(4):511–524.

Lemos, R. T. and Sansó, B. (2006). Spatio-temporal variability of ocean temperature in the Portugal Current System. *Journal of Geophysical Research Oceans*, 111, C4, C04010, 10.1029/2005JC003051.

Lemos, R. T., Sansó, B., and Huertos, M. L. (2007). Spatially varying temperature trends in a Central California estuary. *Journal of Agricultural, Biological, and Environmental Statistics*, 12(3):379–396.

Lima, F. P., Ribeiro, P. A., Queiroz, N., Hawkins, S. J., and Santos, A. M. (2007a). Do distributional shifts of northern and southern species of algae match the warming pattern? *Global Change Biology*, 13(12):2592–2604.

Lima, F. P., Ribeiro, P. A., Queiroz, N., Xavier, R., Tarroso, P., Hawkins, S. J., and Santos, A. M. (2007b). Modelling past and present geographical distribution of the marine gastropod Patella rustica as a tool for exploring responses to environmental change. *Global Change Biology*, 13(10):2065–2077.

R Development Core Team (2005). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Raftery, A. E. and Lewis, S. M. (1992a). Comment: One long run with diagnostics: Implementation strategies for markov chain monte carlo. *Statistical Science*, 7:493–7.

Raftery, A. E. and Lewis, S. M. (1992b). How many iterations in the Gibbs sampler? In Bernardo, J. M., Berger, J. O., Dawid, P., and Smith, A. F. M., editors, *Bayesian Statistics 4*, pages 765–776. Oxford University Press.

Rodionov, S. N. (2006). Use of prewhitening in climate regime shift detection. *Geophysical Research Letters*, 33, L12707, doi:10.1029/2006GL025904.

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Annals of Mathematical Statistics*, 23:470–472.

Santer, B. D., Wigley, T. M. L., Barnett, T. P., and Anyamber, E. (1996). Detection of climate change and attribution of causes. In Houghton, J. T., Meira-Filho, L., Chancellor, B., Kattenberg, A., and Maskell, K., editors, *Climate Change 1995: The Science of Climate Change*, pages 407–444. Cambridge University Press.

Smith, B. J. (2005). *Bayesian Output Analysis Program (BOA) for MCMC*. R package version 1.1.5-2.

Szunyogh, I., Kostelich, E. J., Gyarmati, G., Kalnay, E., Hunt, B. R., Ott, E., Satterfield, E., and Yorke, J. A. (2008). A local ensemble transform Kalman filter data assimilation system for the NCEP global model. *Tellus A*, 60(1):113–130.

von Storch, H. and Zwiers, F. W. (1999). *Statistical Analysis in Climate Research*. Cambridge Univ. Press, NY.

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer Verlag, New York, second edition.

Worley, S. J., Woodruff, S. D., Reynolds, R. W., Lubker, S. J., and Lott, N. (2005). ICOADS release 2.1 data and products. *International Journal of Climatology*, 25(7):823–842.