

A spatiotemporal auto-regressive moving average model for solar radiation

C.A. Glasbey and D.J. Allcroft
Biomathematics and Statistics Scotland
King's Buildings, Edinburgh, EH9 3JZ, Scotland

January 11, 2008

Abstract

To investigate the variability in energy output from a network of photo-voltaic cells, solar radiation was recorded at ten sites every ten minutes in the Pentland Hills to the south of Edinburgh. We identify spatio-temporal auto-regressive moving average (STARMA) models as the most appropriate to address this problem. Although previously considered computationally prohibitive to work with, we show that by approximating using toroidal space and fitting by matching autocorrelations, calculations can be substantially reduced. We find a STAR(1) process with a first-order neighbourhood structure and a Matern noise process to provide an adequate fit to the data, and demonstrate its use in simulating realisations of energy output.

Key words: clearness index, Matern process, spatio-temporal models, Toeplitz block circulant matrices.

1 Introduction

Weather variability in small regions, a few kilometres in size, is important in many hydrological, agricultural and energy contexts. To assess solar energy potential, data were collected from ten radiometers located in the Pentland Hills to the south of Edinburgh (see Figure 1). Solar radiation, averaged over ten minute intervals, was recorded at each site for two years. In this paper we focus on data for the first month, May 1987. Figure 2 shows the data, after transformation to a normalised clearness index, further details of which will be given in §3. There are missing values between sunset and sunrise, when radiation values are zero or close to zero, and occasional other missing values due to equipment problems. The aim of this paper is to build a spatio-temporal model, in order to investigate variability in the energy output from a network of photo-voltaic cells.

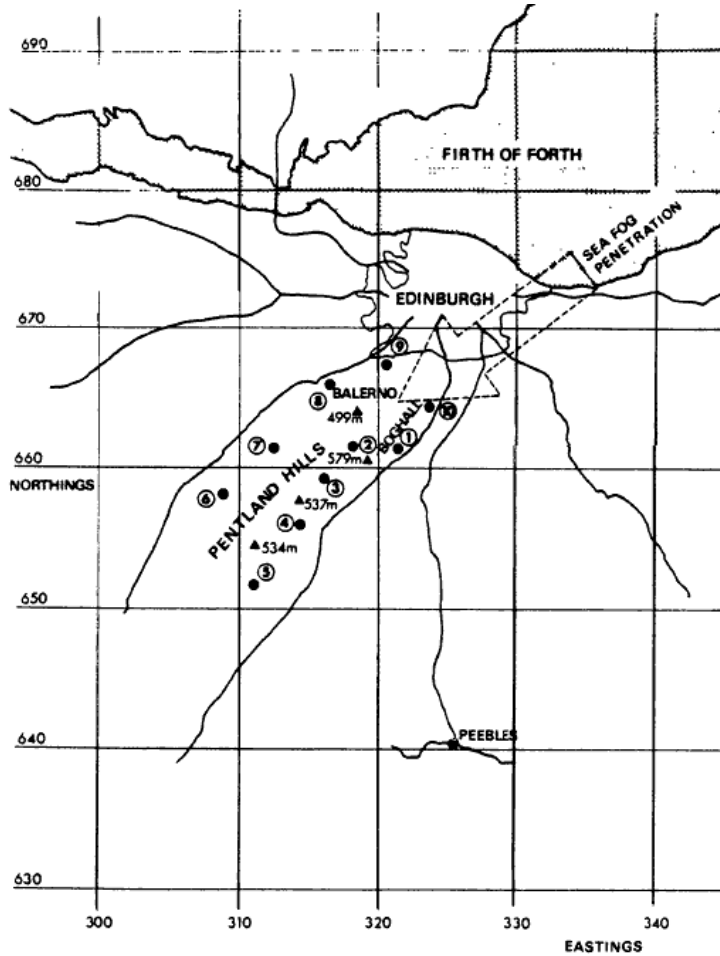


Figure 1: Locations of ten solar radiometers in the Pentland Hills to the south of Edinburgh.

There is currently much interest in spatio-temporal models, with many different approaches dependent on the nature of the available data and modelling objectives. For example, De Luna and Genton (2005) address a situation where data are rich in time but sparse in space, such as we have in Figure 2, by using VAR (vector autoregressive) models. (See Lütkepohl, 1991, for general theory, including VARMA (vector autoregressive moving average) models, which are the multivariate generalisation of ARMA models.) VAR and VARMA models make no assumptions of spatial stationarity, but simply relate the current value at each site to past observations at nearby sites. Although this is a flexible approach, it is only suitable for modelling and simulation at the sites where measurements are taken. Also, it is only tractable for a small number of sites, because otherwise there are too many parameters to be estimated.

As we wish to simulate solar radiation on a network such as a square grid, we turn to STARMA

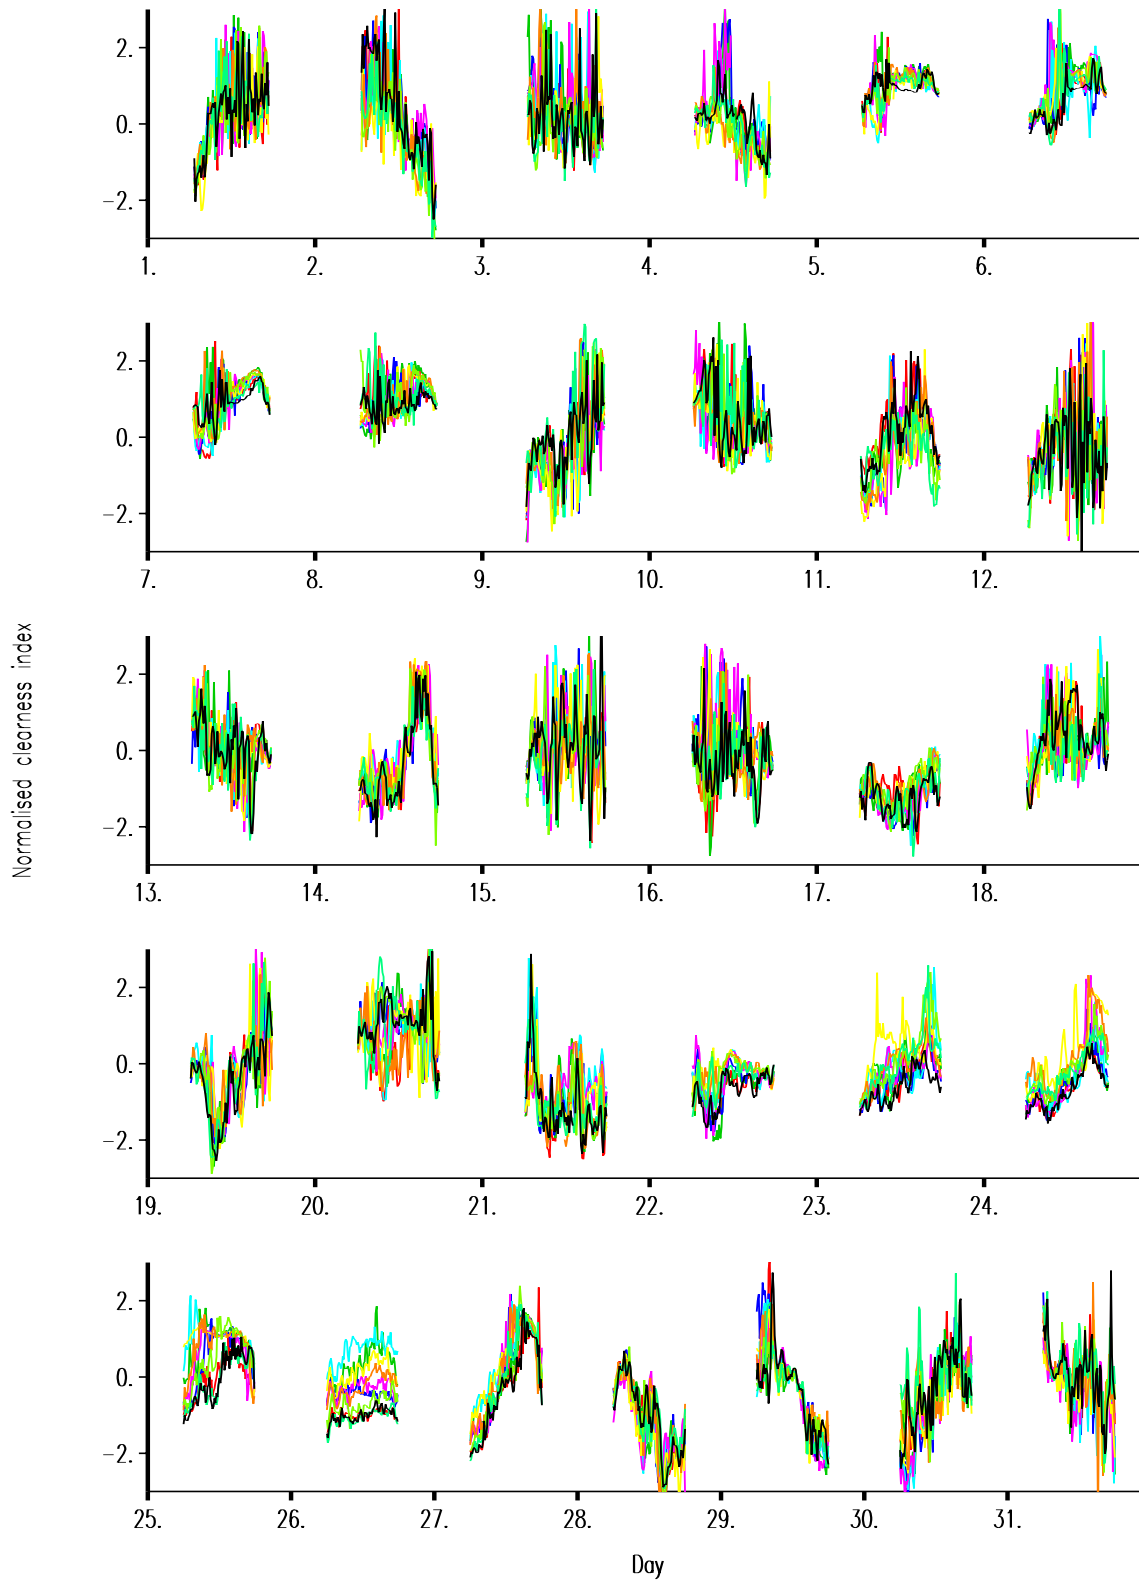


Figure 2: Ten minute averages of solar radiation, after rescaling to a normalised clearness index, plotted against time, for one month's data from all sites (with each site shown as a different shade of grey).

(spatio-temporal auto-regressive moving average) models, which are a special case of VARMA models. STARMA models are defined on a square grid, and parameters are not site specific, but instead simply depend on the vector distances between sites, leading to structured dependence matrices (see, Bennett, 1979). STARMA models were proposed by Pfeifer and Deutsch (1980b). Subsequently, Pfeifer and Deutsch (1980a) addressed issues of model identification, Pfeifer and Deutsch (1981) considered the sampling variance of autocorrelations and Deutsch and Pfeifer (1981) extended the model to allow errors to be correlated. However, since this initial series, there have been surprisingly few papers, either advancing the methodology or using the models, a noteworthy exception being Epperson (2000) who applied them in an ecological context.

In §2 we present new computational methodology for STARMA processes, which we then apply in §3 to model solar radiation data. Finally, in §4 we discuss the results and extensions.

2 STARMA models

The VARMA process of order (p, q) is given by:

$$Y_t = \sum_{r=1}^p \Phi_r Y_{t-r} + \sum_{s=0}^q \Theta_s Z_{t-s} \quad \text{for } t = \dots, -1, 0, 1, \dots \quad (1)$$

Here Y_t is the observed process and Z_t the noise process, both N -vectors, with $Z_t \sim N(0, \Lambda)$ independently at each time t , and Λ is an $N \times N$ positive definite variance matrix. Further, Φ_1, \dots, Φ_p are the autoregressive coefficients and $\Theta_0, \dots, \Theta_q$ are the moving average coefficients, all $N \times N$ matrices, with $\Theta_0 \equiv I$, the identity matrix, for identifiability.

STARMA processes are a special case of VARMA processes, where elements in Y are observations on a square grid. For simplicity, we consider an $n \times n$ grid, and index using a raster scan, so that 2-D grid location (i, j) is represented by element $(ni + j)$ in Y and $N = n^2$. To ensure that the process is spatially homogeneous, an entry at $(ni + j, nk + l)$ in any of the above matrices depends only on the vector distance between 2D locations (i, j) and (k, l) . We further constrain the process to treat row and column distances interchangeably, and therefore be approximately isotropic, and so use Euclidean distances rather than vector distances.

Expected autocovariances of Y can be computed via the Yule-Walker equations (Brockwell and Davis, 1991, p. 420). These are given by:

$$\Gamma_j = \sum_{r=1}^p \Phi_r \Gamma_{j-r} + \sum_{s=j}^q \Theta_s \Lambda \Psi_{s-j}^T \quad \text{for } j = 0, 1, \dots,$$

where $\Gamma_j = E(Y_t Y_{t-j}^T)$, is the $N \times N$ matrix of autocovariances at temporal lag j , and auxiliary matrices, Ψ , are obtained recursively by

$$\Psi_j = \sum_{i=1}^j \Phi_i \Psi_{j-i} + \Theta_j \quad \text{for } j = 0, 1, \dots$$

The first $(p + 1)$ of these matrix equations are solved for $\Gamma_0, \dots, \Gamma_p$, making use of $\Gamma_{-j} \equiv \Gamma_j^T$, and then $\Gamma_{p+1}, \Gamma_{p+2}, \dots$ are obtained recursively. In particular, for a STAR(1) process,

$$Y_t = \Phi Y_{t-1} + Z_t \quad \text{for } t = \dots, -1, 0, 1, \dots,$$

and autocovariances are given by

$$\begin{aligned} \Gamma_0 &= (I - \Phi^2)^{-1} \Lambda \\ \Gamma_j &= \Phi \Gamma_{j-1} = \Phi^j \Gamma_0 \quad \text{for } j = 1, 2, \dots \end{aligned}$$

Unfortunately, even for quite modest sizes of spatial grid, e.g. 64×64 , matrices are of dimension 4096×4096 , and computations become prohibitively expensive. However, by approximating space by a torus, so that row 0 in the grid is adjacent to row $(n - 1)$, and similarly column 0 is adjacent to column $(n - 1)$, all matrices are Toeplitz block circulant (TBC) and the computational burden is considerably reduced using Fourier transforms (see Appendix). Rue and Tjelmeland (2002) used the same approach with Gaussian Markov random fields. For an analysis of the consequences of this approximation in spaces of differing dimensions, see Kent and Mardia (1996). We check its effect in our application in §3. We define a circular distance function d between rows (or columns) i and k :

$$d(i, k) = \min(|i - k|, n - |i - k|),$$

and a toroidal Euclidean distance between 2D locations (i, j) and (k, l) :

$$D(i, j; k, l) = \sqrt{d^2(i, k) + d^2(j, l)}.$$

Terms in matrices only depend on D , so that, for example, a STAR(1) process with spatial dependence of order 2 has

$$\Phi_{ni+j, nk+l} = \begin{cases} \phi_{D^2} & \text{if } D^2(i, j; k, l) \leq 2 \\ 0 & \text{otherwise,} \end{cases}$$

with parameters ϕ_0 , ϕ_1 and ϕ_2 .

The simplest model for the noise process, Z , is $\Lambda \propto I$. However this is unnecessarily constraining, and a more flexible model is obtained by considering spatially correlated noise (see Deutsch and Pfeifer, 1981). We adopt the Matern process, which has been found to be a useful spatial model (see, for example, the published discussion of Diggle et al., 1998). Therefore

$$\Lambda_{ni+j, nk+l} \propto \left(\frac{D(i, j; k, l)}{\beta_1} \right)^{\beta_2} \mathcal{K}_{\beta_2} \left(\frac{D(i, j; k, l)}{\beta_1} \right), \quad (2)$$

where β_1 is a scaling parameter, β_2 is a shape parameter and \mathcal{K}_{β_2} is the modified Bessel function of the third kind of order β_2 (Handcock and Wallis, 1994). Matern processes include, as special cases, the exponential class when $\beta_2 = 1/2$, Whittle's class when $\beta_2 = 1$ and the Gaussian covariance case when $\beta_2 \rightarrow \infty$. We can also allow for independent observation errors, essentially assuming the STARMA process is latent, by replacing Γ_0 by $\Gamma_0 + \sigma^2 I$.

In principle, STARMA model orders can be determined from the patterns of decay or cutoff of correlation and partial correlation functions, as with univariate ARMA models, subject to identifiability problems if $p > 0$ and $q > 0$ (Brockwell and Davis, 1991, p. 431). However, these rapidly become complicated, so model information criteria are easier to use. For parameter estimation, Pfeifer and Deutsch (1980b) discuss the difficulties in using maximum likelihood and propose alternative, conditional likelihood approaches based on least squares. In the solar radiation application, missing values are an extra complication, so in §3 we use simpler, ad-hoc techniques based on matching observed to expected correlations, as did Rue and Tjelmeland (2002) to fit Gaussian Markov random fields. Because autocorrelations can be computed rapidly using properties of TBC matrices, this matching criterion is fast.

STARMA models are particularly suited to simulation, in that once processes Y and Z have been initialised, subsequent time points can be generated directly from (1). The main challenge is to simulate realisations of Y_0 with variance Γ_0 and realisations of Z_t with variance Λ . However, we can again utilise the structure of TBC matrices (see Appendix).

3 Application

The recorded solar radiation values were preprocessed to produce the data plotted in Figure 2. Firstly, diurnal and seasonal trends were removed, by dividing each observed radiation value by the maximum theoretical radiation, based on the elevation of the sun at that position in space and time (Page, 1986). This transformation, which results in a clearness index, was found to be very effective in removing temporal trends in both the mean and variance of solar radiation. Figure 3 shows the histogram of values obtained. The clearness index is close to 100% when the sky is clear and is considerably less when the sky is cloudy. We see that the distribution is bimodal, though not as strongly as for 30-second data, for which Glasbey (2001) developed a new form of nonlinear autoregressive time series model. Unfortunately, there seems to be no way to extend the time series model to encompass spatio-temporal data, and in general there is a paucity of non-Gaussian spatio-temporal models. Therefore, we consider transforming the data to normality. It is straightforward to achieve marginal normality, simply by replacing each observed data value by the corresponding quantile of the standard normal distribution. Of course, this does not guarantee multivariate normality. However, given that the data are not severely non-Gaussian, this is not an unrealistic expectation, and is not contradicted by plots such as Figure 4. Therefore, we will proceed by assuming that the normalised clearness index plotted in Figure 2 is a Gaussian variable.

We estimate the parameters in the STARMA models by numerically minimising the sum of squares of differences between sample and expected correlations:

$$S^2 = \frac{1}{\#} \sum_{i=1}^{10} \sum_{j=1}^{10} \sum_{t \in \mathcal{T}} \left(C_{ijt} - \frac{\Gamma_{0, \Delta_{ij}; t}}{\Gamma_{0,0;0}} \right)^2,$$

where \mathcal{T} is a set of lag times, Δ_{ij} is the Euclidean distance between sites i and j , C_{ijt} is the

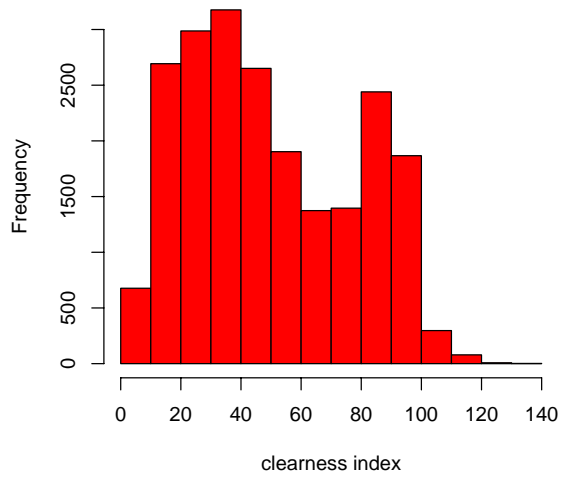


Figure 3: Histogram of values of clearness index, using all data for one month.

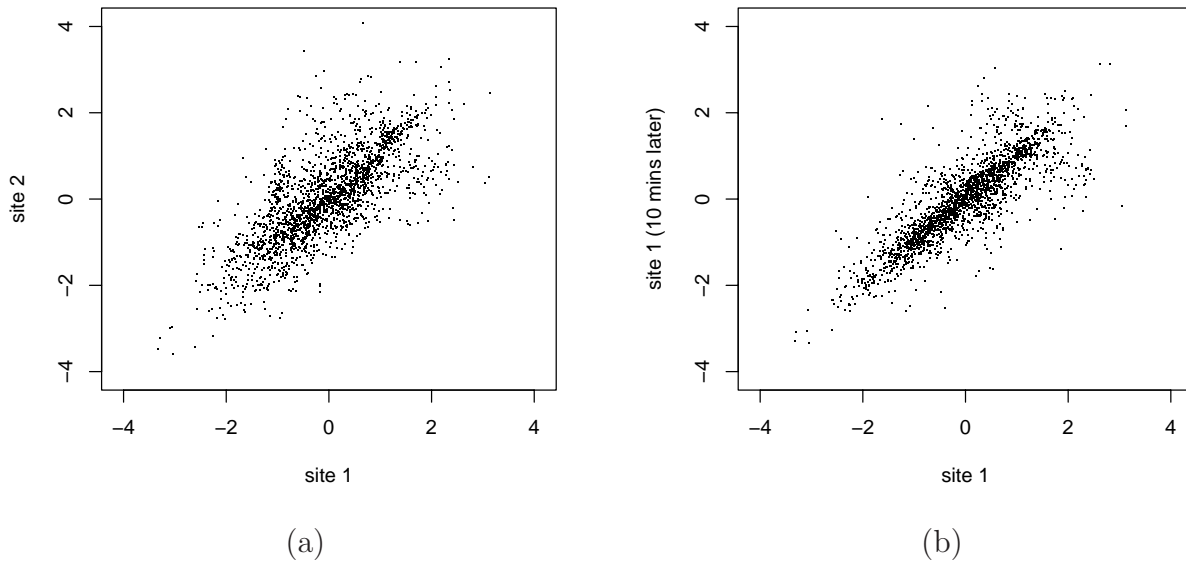


Figure 4: Bivariate plots of clearness index after normalisation: (a) sites 1 and 2 at the same time, (b) site 1 and site 1 ten minutes later.

neighbourhood order	S	$\hat{\phi}_0$	$\hat{\phi}_1$	$\hat{\phi}_2$	$\hat{\beta}_1$	$\hat{\beta}_2$
0	0.0785	0.953			248	0.1819
1	0.0709	0.116	0.210		3115	0.0258
2	0.0704	0.125	0.056	0.151	1587	0.0254

Table 1: *Fits of STAR(1) model for different sizes of neighbourhood.*

sample cross-covariance between sites i and j at temporal lag t :

$$C_{ijt} = \frac{1}{\#} \sum_s X_{i,s} X_{j,s+t} ,$$

$X_{i,s}$ denotes the normalised clearness index at site i time s , and $\#$ denotes the number of terms in each summation. Because X is standardised to zero mean, unit variance, C is also an estimate of the cross-correlation. We assume that the process is approximately isotropic, so that the correlation between sites depends only on the distance between them and we only need consider row-wise distances in Γ .

We initially considered STAR(1) models with a Matern noise process. By restricting time lags to powers of two, i.e. $\mathcal{T} = \{0, 1, 2, 2^2, \dots, 2^6\}$, we considerably speeded-up the fitting of STAR(1) models using the recursive relationship

$$\Phi^{2^t} = \Phi^t \Phi^t \quad \text{for } t = 1, 2, 2^2, \dots .$$

We estimated the STAR processes on a 64×64 grid, with grid points at a spacing of 1km, as the distance between observation sites ranged from 3km to 20km. Table 1 gives results for zeroth-, first- and second-order neighbourhoods, the zeroth-order being particularly quick to fit because in this case space and time are separable, so $\Gamma_j = \phi_0^j \Lambda$. (In general, space-time separability leads to the mathematical convenience of Toeplitz block matrices, but is often incompatible with spatio-temporal data.) We see that a first-order neighbourhood gives a substantial improvement over the zeroth-order one, but increasing to second-order makes little difference to S , with the average difference between sample and expected correlations remaining at 0.07. Nor did a latent STAR process, obtained by modifying the autocorrelations to include an observation error, greatly improve the fit, with $S = 0.0703$. Figure 5 shows excellent agreement between the sample and expected correlations for the STAR(1) model with a first-order neighbourhood. Therefore there seems to be no need to consider larger neighbourhoods or STARMA models of higher order.

To assess the effect of approximating space by a torus, we created a larger torus, on a 128×128 grid, by adding cells around the existing 64×64 grid, and we also computed the autocorrelations of the fitted model directly. A STAR(1) process can be represented as an moving average process of infinite extent:

$$Y_t = \sum_{r=0}^{\infty} \Phi^r Z_{t-r} ,$$

and, in particular, for a first-order neighbourhood, a single element $Y_{ni+j,t} \equiv Y_{(i,j),t}$ is expressible

parameter	true value	100 simulations	
		mean	s.d.
ϕ_0	0.116	0.115	0.0672
ϕ_1	0.210	0.210	0.0132
$\Lambda_{0,1}/\Lambda_{0,0}$	0.379	0.396	0.1062
$\Lambda_{0,20}/\Lambda_{0,0}$	0.258	0.264	0.0655

Table 2: *Results of fitting to 100 simulations of one month's data from STAR(1) model with first-order neighbourhood.*

as

$$Y_{(i,j),t} = \sum_{r=0}^{\infty} \sum_{(k,l) \in \mathcal{S}_{i,j,r}} \Phi_{(i,j),(k,l)}^r Z_{(k,l),t-r},$$

where $\mathcal{S}_{i,j,r}$ denotes the set $\{(k,l) : |k-i| + |j-l| \leq r\}$. Therefore, as Z_t and $Z_{t'}$ are uncorrelated if $t \neq t'$, autocovariances can be computed using

$$\text{cov}(Y_{(i,j),t}, Y_{(i',j'),t'}) = \sum_{r=0}^{\infty} \sum_{(k,l) \in \mathcal{S}_{i,j,r}} \sum_{(k',l') \in \mathcal{S}_{i',j',r+t'-t}} \Phi_{(i,j),(k,l)}^r \Phi_{(i',j'),(k',l')}^{r+t'-t} \Lambda_{(k,l),(k',l')} \quad \text{for } t' \geq t,$$

where elements in Φ^r are computed recursively from those in Φ^{r-1}

$$\Phi_{(i,j),(k,l)}^r = \sum_{(i',j') \in \mathcal{S}_{i,j,1}} \Phi_{(i,j),(i',j')} \Phi_{(i',j'),(k,l)}^{r-1}.$$

Reassuringly, we found that differences in autocorrelations between the 64×64 torus and exact calculations were never greater than 3×10^{-5} , which occurred at the maximum distance considered of 20km, and this reduced to 10^{-7} when using a 128×128 torus. We conclude that results are insensitive to the torus approximation.

We used the methods described in the Appendix to simulate realisations of the fitted STAR(1) model. To assess the estimation method for bias and precision, we first generated 100 independent replicates of one month of data matching the original data. We did this by simulating sequences of 31 days on the 64×64 grid, then extracting data for the 10 sites, restricted to times for which radiation was recorded. Figure 6 shows one realisation at a single instant in time, after transforming back to the scale of clearness index. A pattern reminiscent of clouds is apparent. We estimated parameters as for the original data, and Table 2 summarises the results. We found the β -parametrisation of the Matern process (2) to be unstable, with $\hat{\beta}_1$ ranging from 190 to 9×10^7 and $\hat{\beta}_2$ from 7×10^{-11} to 0.10, whereas autocorrelations varied far less, so we report these at lags 1 and 20 in the table. There is no evidence of bias in any of the parameter estimates. Diggle and Ribeiro (2007, pp. 51-53) show that the main effect of the shape parameter is to generate a smooth process if $\beta_2 > 1$, which is otherwise rough, so we could alternatively have set β_2 to a value less than unity rather than attempting to estimate it.

To address the issue of variability in the energy output from a square grid of photo-voltaic cells, we computed the average value of the clearness index for the 10×10 km square shown in the centre of Figure 6, given that energy output will be a function of this index. For illustration, we

simulated 999 independent realisations of 12 hours of data, starting from a low level of clearness index, and generated time series of average clearness index within the square. Figure 7 shows a single realisation together with median and 2.5% and 97.5% ranges of values, which give an indication of the trends and patterns of fluctuation in energy output.

We also simulated a longer, stationary time series, of length 10000. The marginal distribution was approximately Gaussian, and sample autocorrelations are shown in Figure 8. A univariate ARMA(1,1) process was found to give an excellent fit to this time series, with $\hat{\phi} = 0.942$, $\hat{\theta} = 0.449$, as can be seen from the fitted autocorrelations also shown in Figure 8. Increasing the order of ARMA model did not significantly improve the fit. We conclude that, to study stationary behaviour of output from a grid of photo-voltaic cells, such as a 10×10 km square, it is sufficient to consider a univariate ARMA(1,1) process. Of course, for areas of different shape and size, model orders and parameter values will need to be modified.

4 Discussion

We have seen that STARMA models offer a convenient framework for modelling spatio-temporal data on a grid. The formulation is particularly well suited to forward simulation in time. By approximating by a toroidal space and fitting by matching autocorrelations, the calculations can be reduced substantially and this makes the models far more tractable. In the application, we found a STAR(1) process with a first-order neighbourhood structure and a Matern noise process to provide an adequate fit. Further, we demonstrated the usefulness of the model in simulating spatio-temporal realisations of solar radiation in order to assess variability in energy output. However, the fitting procedure was ad hoc. In further work it would be useful to investigate the efficiency of the method in comparison with alternatives that are more computationally intensive, but have better theoretical bases.

Gaussian Markov random fields (GMRFs) offer an alternative modelling framework for spatio-temporal data, which we used in earlier work on rainfall disaggregation (Allcroft and Glasbey, 2003). GMRFs were particularly suited to disaggregation, as this necessitates simulating from conditional distributions. In the solar radiation application, the objective is to simulate forwards in time, which is problematic using GMRFs. However, despite using different models in the two papers, there are similarities in model fitting strategy. In both cases we use a toroidal approximation and hence make use of the savings in computational time afforded by TBC matrices, the ideas having first been outlined in the GMRF context by Rue and Tjelmeland (2002).

Partially ordered Markov models (POMMs) offer yet another alternative, in some sense intermediate between GMRFs and STARMA processes. Cressie and Davidson (1998) and Davidson et al. (1999) give the main details and use the models in a spatial context for generating textures in images, though see also Baddeley et al. (2002) for some useful comments on partial ordering. Although spatial data lack the natural ordering of temporal data, POMMs exploit at least a partial ordering. A spatio-temporal POMM can be specified with a partial ordering such that conditional distributions depend only on observations at previous time points.

However, this form of POMM is equivalent to a STAR model. POMMs are special cases of MRFs, with a POMM neighbourhood structure determining the neighbourhood structure of the corresponding MRF. However in general, given the neighbourhood size, the POMM will have less parameters and from some initial attempts to fit these models to the rainfall data in Allcroft and Glasbey (2003), it was found that the extra flexibility of the MRF was important.

STARMA, GMRF and POMM models are all defined on a grid that is discrete in both space and time, whereas solar radiation, in common with most other real-world processes, exist in continuous space and time. In particular, the choice of a grid with 1km spacing is somewhat arbitrary. However, if a model had been identified on a finer grid, say with a $\frac{1}{2}$ km spacing, then the implied process on the original, coarser grid would not have been exactly a STARMA, GMRF or POMM model of any order, although it could be closely approximated. If it is important to have a representation in continuous space and time, then the simplest option is to model the spatio-temporal autocorrelation function directly, but then forward simulation would be far more problematic.

Acknowledgements

The work was supported by funds from the Scottish Government. Data collection was funded by the European Union, under contract EN35-0049-D(B)/115.

References

- Allcroft, D. J. and Glasbey, C. A. (2003). A latent Gaussian Markov random field model for spatio-temporal rainfall disaggregation. *Applied Statistics*, 52:487–498.
- Baddeley, A., Nair, G., and Cressie, N. (2002). Directed Markov point processes: characterisation and construction. Technical Report 2002/14, Department of Mathematics and Statistics, University of Western Australia.
- Bennett, R. J. (1979). *Spatial Time Series: Analysis–Forecasting–Control*. Pion, London.
- Brockwell, P. J. and Davis, R. A. (1991). *Time Series: Theory and Methods*. Springer-Verlag, New York, second edition.
- Cressie, N. and Davidson, J. L. (1998). Image analysis with partially ordered Markov models. *Computational Statistics & Data Analysis*, 29:1–26.
- Davidson, J. L., Cressie, N., and Hua, X. (1999). Texture synthesis and pattern recognition for partially ordered Markov models. *Pattern Recognition*, 32:1475–1505.
- De Luna, X. and Genton, M. G. (2005). Predictive spatio-temporal models for spatially sparse environmental data. *Statistica Sinica*, 15:547–568.

- Deutsch, S. J. and Pfeifer, P. E. (1981). Space-time ARMA modeling with contemporaneously correlated innovations. *Technometrics*, 23:401–409.
- Diggle, P. J. and Ribeiro, P. J. (2007). *Model-based Geostatistics*. Springer, New York.
- Diggle, P. J., Tawn, J. A., and Moyeed, R. A. (1998). Model-based geostatistics (with discussion). *Applied Statistics*, 47:299–350.
- Epperson, B. K. (2000). Spatial and space-time correlations in ecological models. *Ecological Modelling*, 132:63–76.
- Glasbey, C. A. (2001). Nonlinear autoregressive time series with multivariate Gaussian mixtures as marginal distributions. *Applied Statistics*, 50:143–154.
- Handcock, M. S. and Wallis, J. R. (1994). An approach to statistical spatial-temporal modeling of meteorological fields (with discussion). *Journal of the American Statistical Association*, 89:368–390.
- IMSL (2002). *IMSL. Math/Library, Stat/Library*. Visual Numerics, Inc., Houston, Texas.
- Kent, J. T. and Mardia, K. V. (1996). Spectral and circulant approximations to the likelihood for stationary Gaussian random fields. *Journal of Statistical Planning and Inference*, 50:379–394.
- Lütkepohl, H. (1991). *Introduction to Multiple Time Series Analysis*. Springer-Verlag, Berlin.
- Page, J. K. (1986). Prediction of solar radiation on inclined surfaces. In *Solar Energy R & D in the European Community: Series F, Volume 3 – Solar Radiation Data*. D. Reidel Publishing Company, Dordrecht.
- Pfeifer, P. E. and Deutsch, S. J. (1980a). Identification and interpretation of first order space-time ARMA models. *Technometrics*, 22:397–408.
- Pfeifer, P. E. and Deutsch, S. J. (1980b). A three-stage iterative procedure for space-time modeling. *Technometrics*, 22:35–47.
- Pfeifer, P. E. and Deutsch, S. J. (1981). Variance of the sample space-time autocorrelation function. *Journal of the Royal Statistical Society, Series B*, 43:28–33.
- Rue, H. and Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields. *Scandinavian Journal of Statistics*, 29:31–49.

Appendix: Toeplitz block circulant matrices

By approximating space by a torus, all matrices are Toeplitz block circulant (TBC). Toeplitz circulant matrices have the form

$$A = \begin{pmatrix} a_0 & a_1 & a_2 & \cdots & a_{n-1} \\ a_{n-1} & a_0 & a_1 & \cdots & a_{n-2} \\ a_{n-2} & a_{n-1} & a_0 & \cdots & a_{n-3} \\ \vdots & \vdots & \vdots & & \vdots \\ a_1 & a_2 & a_3 & \cdots & a_0 \end{pmatrix}$$

(Brockwell and Davis, 1991, pp. 133-136). For TBC matrices, each a_k is itself an $n \times n$ Toeplitz circulant matrix. The structure is such that each row of A contains the same elements, but in a different order. So, all information is contained in the first row of the $n^2 \times n^2$ matrix, and can be reformatted as an $n \times n$ matrix, and denotes terms linking the first-indexed location in the torus, at $(0, 0)$, with all other locations.

The product of two TBC matrices is also TBC, and can be computed using only the elements in the first row:

$$(AB)_{00,kl} = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} A_{00,ij} B_{00,\text{mod}_n(k-i),\text{mod}_n(l-j)} \quad \text{for } k, l = 0, \dots, n-1.$$

Further, the inverse of a TBC matrix, A , is TBC and can be computed efficiently. The first step is to apply a two-dimensional Fourier transform to its first row after reformatting as an $n \times n$ matrix (denoted A_{00}), to obtain a matrix of eigenvalues, A^* :

$$A_{ij}^* = \sum_{k=0}^{n-1} \sum_{l=0}^{n-1} A_{00,kl} \exp \left\{ -2\pi\iota \left(\frac{ik}{n} + \frac{jl}{n} \right) \right\} \quad \text{for } i, j = 0, \dots, n-1,$$

where $\iota \equiv \sqrt{-1}$. The second step is an inverse Fourier transform of the inverse eigenvalues to obtain the elements of the first row of the inverse matrix:

$$A_{00,kl}^{-1} = \frac{1}{n^2} \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} \frac{1}{A_{ij}^*} \exp \left\{ 2\pi\iota \left(\frac{ik}{n} + \frac{jl}{n} \right) \right\} \quad \text{for } k, l = 0, \dots, n-1.$$

The Fourier transforms are computationally fast; if the dimensions n are products of small primes, then computational effort is proportional to $n^2 \log n$ (see, for example, routines DFFT2D and DFFT2B in IMSL, 2002).

Finally, TBC structure simplifies simulation. For example, to simulate a grid of spatial data Z with TBC variance matrix Λ , we first compute the Fourier coefficients

$$\Lambda_{jk}^* = \sum_{l=0}^{n-1} \sum_{m=0}^{n-1} \Lambda_{00,lm} \exp \left\{ -2\pi\iota \left(\frac{jl}{n} + \frac{km}{n} \right) \right\} \quad \text{for } j, k = 0, \dots, n-1.$$

We then simulate complex normal deviates, where z_{jk}^* has mean zero, variance Λ_{jk}^* in both real and imaginary components. Lastly we perform an inverse 2D Fourier transform:

$$Z_{jk} = \sum_{l=0}^{n-1} \sum_{m=0}^{n-1} z_{lm}^* \exp \left\{ 2\pi\iota \left(\frac{jl}{n} + \frac{km}{n} \right) \right\} \quad \text{for } j, k = 0, \dots, n-1,$$

resulting in an array Z_{jk} of simulated data with the required covariances.




Figure 5: Cross-correlations between normalised clearness indexes at a range of spatial and temporal lags: (×) sample values, (—) fitted STAR(1) model with a first-order neighbourhood.



Figure 6: Simulation of clearness index on 64×64 grid using fitted STAR(1) model, with smaller values shown as lighter shades of grey. Also shown is a 10×10 km region of interest for a network of photo-voltaic cells.




Figure 7: Simulation of clearness index averaged over a 10×10 km square starting from a low value, plotted against time (dark line), together with median and 2.5% and 97.5% ranges of values (pale lines).




Figure 8: Autocorrelations from time series of length 10000 of average of 10×10 km square of the clearness index: \circ sample values, — fit of ARMA(1,1) model.