RESEARCH ARTICLE

Comparison of three dimensional profiles over time

Margaret R. Donald^{a*}, Chris Strickland^a, Clair L. Alston^a, Rick Young^b and Kerrie L. Mengersen^a

^a*School of Mathematical Sciences, Queensland University of Technology, GPO Box 2434, Brisbane, QLD 4001, Australia;*

^b*Tamworth Agricultural Institute, Industry & Investment NSW, 4 Marsden Park Road, Calala NSW 2340, Australia*

(Received 00 Month 200x; in final form 00 Month 200x)

We describe an analysis for data collected on a three-dimensional spatial lattice with treatments applied at the horizontal lattice points. Spatial correlation is accounted for using a conditional autoregressive (CAR) model. Observations are defined as neighbours only if they are at the same depth. This allows the corresponding variance components to vary by depth. We use Markov Chain Monte Carlo (MCMC) with block updating, together with Krylov subspace methods for efficient estimation of the model. The method is applicable to both regular and irregular horizontal lattices and hence to data collected at any set of horizontal sites for a set of depths or heights, for example, water column or soil profile data. The model for the three-dimensional data is applied to agricultural trial data for five separate days taken roughly six months apart, in order to determine possible relationships over time. The purpose of the trial is to determine a form of cropping that leads to less moist soils in the root zone and beyond. We estimate moisture for each date, depth and treatment accounting for spatial correlation and determine relationships of these and other parameters over time.

Keywords: Bayesian; Conditional Autoregressive (CAR) models; depth profiles; Field trial; Linear spline; Markov Chain Monte-Carlo (MCMC); Gaussian Markov random field (GMRF); Spatial autocorrelation; Variance components.

*Corresponding author. Email: merricks.merricks@gmail.com

1. Introduction

Despite numerous papers examining crop rotations and field experiments conducted over lengthy periods of time, it remains a difficult problem to analyse such data satisfactorily, and in this case, the problem is not that of a response measured over a surface, but that of a response measured over three dimensions of the field. We describe field trial data from a long term crop rotation trial, conducted to determine a cropping system which would maximise the use of stored water in the soil, and minimise the risk of water leakage leaching the soil of its salts and endangering long-term agriculture on the Liverpool Plains in New South Wales, Australia.

These data pose several problems: how to describe the treatment effect, how to account for spatial autocorrelation, how to account for spatial correlation over depths, and what might be an appropriate model over time.

The layered CAR model, a three dimensional variant of the CAR model [7, 9], was shown to well describe the data for a single day [14]. We use an identical layered CAR model for each date to analyse five sets of soil moisture measurements taken six months apart over a two-year period. The choice of CAR models is further discussed in Section 5.

The model for the treatment effect assumes that each treatment determines a depth profile curve for each date. Treatment effects are modelled as continuous curves along the depth dimension, with different curves for each date and treatment. Linear splines are used to model treatment effect over depth allowing trend comparisons over segments of the curve for the different treatments and dates.

In addition to the agricultural case study, we present a computationally efficient Gibbs sampling algorithm for fitting the layered CAR model with a regression component.

Section 2 describes the data used in the case-study. Section 3 describes the model and its computational framework. Section 3 also describes the methods used for comparisons of contrasts between and within dates. Section 4 provides the results of the case-study. Section 5 provides a discussion of the methods and the results. Section 8 outlines a Gibbs sampler for the proposed model.

2. Case Study

The four dimensional data used in this case-study consist of moisture observations taken at 108 surface treatment sites, 15 depths and over 5 different dates during a two-year period. The 108 measurement sites are arranged as 6 rows with 18 columns per row. Hence, data at each time point consist of 1620 measurements at 108 sites over 15 depths.

The purpose of the field trial is to determine a cropping system which leads to lower moisture values in the soils, in order to minimise the risk of deep drainage. More complete details of the trial may be found in Ringrose-Voase et al. [33]. Nine treatments are considered. These fall into three groups, long fallow cropping, response cropping and pasture treatments.

The primary question of interest to crop scientists is whether response cropping gives lower moisture values both at the intermediate and greater depths, in comparison with long fallow, and whether this is sustained over different stages of the cropping cycle. Subsidiary questions addressed here are the comparison of cropping treatments with pasture treatments, and the comparison of the lucerne pasture mixtures with the native grass pasture.

The concern in this paper is to establish whether the various components of the model vary from date to date, and to determine the cropping system best suited to the land.

The treatments are

- (1) Treatments 1-3: Long fallow wheat/sorghum rotation, where one wheat and one sorghum crop are grown in three years with an intervening 10-14 month fallow period. The 3 treatments are each of 3 phases of the long fallow 3 strip system.
- (2) Treatment 4: Continuous cropping in winter with wheat and barley grown alternately.
- (3) Treatments 5 and 6: Response cropping, where an appropriate crop (either a winter or a summer crop) is planted when the depth of moist soil exceeded a predetermined level.
- (4) Treatments 7-9: Perennial pastures. The three treatments are lucerne (a deep rooted perennial forage legume with high water use potential), lucerne

grown with a winter growing perennial grass, and a mixture of winter and summer growing perennial grasses.

The dates of the observations are almost equally spaced over two years (July 23, 1997, December 4, 1997, April 28, 1998, September 23, 1998, and February 25, 1999).

3. Methods

3.1 Model

The model describes data collected over the three spatial dimensions, in particular, over a three dimensional lattice, with the response variable arising from an experimental treatment applied at lattice points on the horizontal plane. The model consists of a regression or fixed effect component, a spatial component and an irregular component or residual error. For the spatial component we consider both non-stationary and stationary spatial processes. The regression component is specified via the treatments (determined by the horizontal planar spatial locations), and a set of basis functions over the third spatial dimension, depth. As in the case of two-dimensional spatial models, we enumerate the horizontal spatial locations in a particular order which determines the spatial neighbourhood matrix. To simplify the definition of the matrices associated with the variances and precisions, the $(n \times 1)$ vector of observations, $\mathbf{y}$, is arranged as $D \times S$ observations, where D is the number of lattice points in the depth dimension, and S the number of spatial locations in the horizontal plane.

The model for $\mathbf{y}$ is as follows

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\psi} + \boldsymbol{\epsilon}, \quad (1)$$

where $\mathbf{X}$ is an $(n \times p)$ design matrix, $\boldsymbol{\beta}$ is a $(p \times 1)$ vector of regression coefficients, $\boldsymbol{\psi}$ is an $(n \times 1)$ vector that models spatial correlation at each depth and $\boldsymbol{\epsilon}$ is an $(n \times 1)$ residual vector that is homogeneous within each depth. The design matrix, $\mathbf{X}$, models the treatment effects as continuous functions of depth for each treatment. The spatial covariance is modelled using a Gaussian Markov random field (GMRF). Stationary and non-stationary covariance structures are considered.

A proper conditional autoregressive (CAR) prior [15] is used for the stationary case, while an intrinsic CAR prior [9, 35] is used in the non-stationary case. The spatial variation for $\boldsymbol{\psi}$ is captured either through a proper prior on $\boldsymbol{\psi}$ in the stationary case such that

$$\boldsymbol{\psi} \sim N\left(\mathbf{0}, \boldsymbol{\Omega}(\rho, \boldsymbol{\tau})^{-1}\right),$$

or in the non-stationary case, through an improper prior

$$\boldsymbol{\psi} \sim N\left(\mathbf{0}, \boldsymbol{\Omega}(\boldsymbol{\tau})^{-1}\right).$$

Points on the lattice are defined as neighbours only if they lie in the same horizontal layer. The precision matrices, $\boldsymbol{\Omega}(\rho, \boldsymbol{\tau})$ and $\boldsymbol{\Omega}(\boldsymbol{\tau})$, are $(n \times n)$ block diagonal matrices that depend on the horizontal neighbourhood structure, the $(D \times 1)$ vector of scaling coefficients, $\boldsymbol{\tau}^2$, and, in the stationary case, a spatial dependence parameter ρ , where $|\rho| < 1$. In the non-stationary case ρ is not required. The block diagonal structure permits D separate scaling coefficients, $\boldsymbol{\tau}^2$, that model differing variances at each depth for the spatial components.

The error, $\boldsymbol{\epsilon}$, is an $n \times 1$ vector, that is defined such that

$$\boldsymbol{\epsilon} \sim N(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\sigma})),$$

where $\boldsymbol{\Sigma}(\boldsymbol{\sigma})$ is an $(n \times n)$ diagonal covariance matrix that is a function of a $(D \times 1)$ vector $\boldsymbol{\sigma}^2$ that allows heterogeneity across depths in the non-spatial random component. The variance, $\boldsymbol{\Sigma}$, is defined as

$$\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2 \mathbf{I}, \sigma_2^2 \mathbf{I}, \dots, \sigma_D^2 \mathbf{I}),$$

where $\mathbf{I}$ is the $S \times S$ identity matrix. This structure arises from the ordering of $\mathbf{y}$ by depth and then by spatial site. The residuals are modelled as having differing variances at each depth.

Two variables, depth and treatment, are used to describe the fixed effects. Here, depth is treated as a continuous variable, and a set of basis functions is formed from it in order to fit splines. For the case-study, the basis functions are linear splines, but other basis functions may be used [25]. Treatment is a categorical variable,

with T levels. The design matrix, $\mathbf{X}$ may be expressed as

$$\mathbf{X} = \mathbf{A} \otimes \mathbf{B},$$

where $\mathbf{A}$ is a $D \times k$ matrix of k depth basis functions, $\otimes$ is the Kronecker product and $\mathbf{B}$ is an $S \times T$ matrix that matches the S horizontal sites with the appropriate set of T dummy variables for the site treatments. The Kronecker product gives $\mathbf{X}$ as an $n \times p$ matrix with $n = D \times S$ and $p = k \times T$.

The linear splines are defined as $z_k(d) = (d - \kappa_k)_+$, where

$$\begin{aligned} (d - \kappa_k)_+ &= d - \kappa_k, \quad d \geq \kappa_k, \\ &= 0, \quad d < \kappa_k, \end{aligned}$$

for some knot sequence $\kappa_1, \dots, \kappa_{k-2}$, and $d = 1, 2, \dots, D$. The basis functions in $\mathbf{A}$ for each d are $[1, d, z_1(d), z_2(d), \dots, z_{k-2}(d)]$ [25]. For the linear splines used in the case-study of this paper, the number of basis functions, k , is the number of internal knots plus 2.

For the stationary CAR prior the precision matrix is

$$\mathbf{\Omega}(\rho, \boldsymbol{\tau}) = \text{block diagonal}(\mathbf{Q}/\tau_1^2, \mathbf{Q}/\tau_2^2, \dots, \mathbf{Q}/\tau_D^2),$$

with $\boldsymbol{\tau}^2$ an $n \times 1$ vector of scaling coefficients permitting different variances at the D different depths, and $\mathbf{Q}$ an $S \times S$ first order neighbourhood precision matrix common to each depth layer, and

$$\mathbf{Q} = (\mathbf{M} - \rho \mathbf{W}).$$

The neighbourhood matrix, $\mathbf{W}$, is defined such that

$$w_{ij} = \begin{cases} 0 & i = j, \\ -1 & i \sim j, \text{ } (i, j \text{ are neighbours}), \\ 0 & \text{otherwise,} \end{cases}$$

and $\mathbf{M}$ is given by

$$\mathbf{M} = \text{diag}(n_1, n_2, \dots, n_S),$$

where n_i is the number of neighbours of site i . See Gelfand and Vounatsou [15].

In the depth layered scheme used here the non-stationary CAR prior is defined as

$$\mathbf{\Omega}(\boldsymbol{\tau}) = \text{block diagonal}(\mathbf{R}/\tau_1^2, \mathbf{R}/\tau_2^2, \dots, \mathbf{R}/\tau_D^2),$$

with $\mathbf{R}$ an $S \times S$ first order neighbourhood precision matrix whose elements r_{ij} are specified by

$$r_{ij} = \begin{cases} n_i & i = j, \\ -1 & i \sim j, \\ 0 & \text{otherwise,} \end{cases}$$

where n_i is the number of neighbours for site i [35].

3.2 Computation

Computation is performed using a general-purpose MCMC software framework currently under development, which allows block updating of parameters. Programming is in Python and uses the Fortran and C libraries, LAPACK, BLAS, SciPy of Anderson et al. [2], Blackford et al. [10], and NumPy Community [26] respectively. Model parameters are partitioned into five blocks, $(\boldsymbol{\psi}, \boldsymbol{\tau}, \boldsymbol{\beta}, \boldsymbol{\sigma}, \rho)$, each jointly sampled. Closed form samplers are used for all model parameters except ρ where a Metropolis Hastings sampler is used. Block updating is found to be more efficient by various authors, see, for example, Chib and Carlin [13], Pitt and Shephard [30]. Lui et al. [21] show theoretically that jointly sampling parameters in a Gibbs scheme leads to a reduction in correlation in the associated Markov chain in comparison with the individual sampling of parameters. Block updating typically means that MCMC chains converge faster.

The conditional autogressive models have a sparse precision matrix defined by the adjacency matrix. The sparse matrix representation used here is the compressed sparse row format described by Saad [36]. Krylov subspace methods are used for updating [37, 40].

Further computing details are given in Section 8, where block updating equations are given for the posterior probabilities.

3.3 Fixed effects

Three models are considered for the regression component: a three-knot linear spline with knots at depth indices 4, 7 & 10, a five-knot linear spline with equally spaced knots at 3.33, 5.67, 8, 10.33 and 12.67, and a saturated model of 135 terms which fits a constant for each treatment by depth.

The linear splines allow discussion of trends across various depth segments. The five-knot linear spline was the initial choice. However, with six segments defined over 15 depth points, linear trends may not be seen because of the limited number (3) of different depths within a segment, so a three-knot model was also considered. For comparison, a saturated model of treatments by depths or $9 \times 15 = 135$ parameters was also fitted.

Smooth continuous curves may be fitted using the generalised additive (GAM) framework of Hastie and Tibshirani [17], or the random walk (order 2) (RW2) smoothing of INLA [23], the RW2 penalised splines of BayesX [3, 4] which are described more fully in Brezger and Lang [11], Lang and Brezger [19] and Kneib and Fahrmeir [18]. Such frameworks seem unnecessarily complicated for the problem here. (For example, a naive use of BayesX, gave a default 20 knots across the 15 depth values.) Additionally, the use of linear splines allows comparisons of trends over the linear segments of the curves.

Choice of the regression component and final model is made using the Deviance Information Criterion (*DIC*) [39], an information criterion based on the Deviance and adjusted by an estimated number of parameters. The results of these comparisons are reinforced by the curves of the posterior deviance distributions [1]. See Figure 1.

3.4 Contrast and parameter comparisons

Output for each MCMC simulation after burnin was kept for all model estimates. This permitted post-hoc comparisons for any desired function both within and across the measurement dates. Contrasts of interest are (1) Average Long fallow cropping minus average response cropping, (2) Average cropping minus average pastures, and (3) Average lucernes minus native perennial pastures. These contrasts are calculated for both the slope of the line segment from 200 cm to 300 cm, and

for the moisture estimates at each point in the depth profile.

Contrasts and slopes are compared across all combinations of dates, giving 10 comparisons for each estimate. Comparisons within a date are formed by pairing the estimates from the same iteration. However, the estimates from each date's model are independent. Hence, the across date comparisons are formed after randomising the iterates.

We compare the variance components of the model in the same manner. For the random spatial components a visual comparison only is made, using the 95% credible intervals for ψ for each site and depth. For depths from 140 cm and onward, these credible intervals largely overlap.

3.5 *Convergence and accuracy*

Geweke statistics [16] were used to assess convergence for all contrasts and comparisons discussed, both for treatments, and variance components. The accuracy of the credible intervals was assessed using Raftery-Lewis statistics [32].

4. Results

4.1 *Model choice*

The DIC (Table 1) indicates that the three-knot linear spline model with the stationary CAR prior is a better model than the three-knot model with the non-stationary CAR, and a more appropriate model than the stationary CAR prior saturated model or the five-knot linear spline on almost all of the five dates. On the date (date 4, September 23, 1998) when the five-knot model is found to be the best, the three-knot model is virtually equivalent. Clearly, tracking treatment effects at depths where they do not exist leads to a poorer fit. However, it seems likely that the improvement observed with the five-knot model on September 23, 1998 represents a slightly better fit at the shallower depths for that date. The three-knot linear spline fits the bulk of the data well, but may be a poorer model for some dates at the shallower depths covered by the first linear segment.

The Deviance curves for the three-knot linear spline, five-knot linear spline and saturated model show the superiority of the five-knot linear spline model for date 4 (Figure 1). Plots of deviance curves for all models and dates show the saturated

model generally to be the poorest of the fitted models.

A useful byproduct of the DIC calculation is the calculation of pD , the effective number of parameters. The saturated model contains 135 parameters and the variance components consist of 31 parameters, but as can be seen from Table 1, pD is approximately equal $5/8$ of the degrees of freedom available on any date. This proportion indicates that many of the spatial residual components might well be considered to be outliers of the CAR normal models [39].

4.2 Convergence and Accuracy Results

Convergence was satisfactory for all parameters discussed and these were found to be accurate to the levels displayed in the paper, using Raftery-Lewis tests [31, 32]. Dependence factors [31] were less than five for all parameters. Geweke [16] z statistics were less than $|2|$. These indicate both satisfactory convergence and that the reported quantiles have the accuracy indicated.

4.3 Variance components

Figures 2 and 3 show the square roots of the non-spatial and the spatial variance component parameters, σ^2 and τ^2 , and indicate that they mirror each other at the various depths and over the dates. Figure 2 illustrates the need for a non-spatial variance component for each depth, but that these components may be constant over dates. The comparisons to see whether the non-spatial variance component for each depth differs across dates, show that all 100 of the possible comparisons across dates for depths from 120 cm to 300 cm have 95% credible intervals which include zero. For depths from 20 cm to 100 cm (50 comparisons) just 9 differences have credible intervals not inclusive of zero, and these all involve comparisons with date 4. These intervals indicate that the non-spatial variance component, σ^2 , varies by depth but not by date.

Figure 3 shows τ varying by depth, but being approximately constant across dates from depths 120 cm to 300 cm. Comparisons across dates show just 3 observed differences whose 95% credible interval fails to include zero from a possible 100. There are apparently some differences across dates at the shallower depths, with 25 of the 50 possible comparisons showing differences for depths from 20 cm to 100 cm,

and these are generally differences with the τ values for date 4. The spatial variance components vary by depth, but not by generally by date. This is particularly true for depths from 120 cm to 300 cm.

The variance component graphs (Figures 2, 3) show very much lower variability in the mid-depth range. Date 4 (September 23, 1998) shows considerably smaller variances for the shallower depths than those for the other dates for both the spatial and non-spatial variance components.

Tables 2 and 3 show values and comparisons for the parameter ρ . Just one of the possible 10 comparisons across dates has a 95% credible interval which did not include zero. ρ appears to be effectively the same across dates.

4.4 Depth segments and dates

The three-knot spline model consists of four linear segments for each treatment. Table 4 shows the 95% credible intervals for the slope of the linear segment at the greatest depth (from 200 cm to 300 cm). Almost all treatments show no trend in this segment. (The exceptions are treatment 8, a lucerne mixture treatment which shows decreasing moisture in this line segment, and treatment 2 which on two of the five dates shows increasing moisture.) In general, the last linear segment (from 200 cm - 300 cm) is constant for all treatments over all dates. Hence, from about 200 cm depth and deeper, the treatments would appear to no longer affect the moisture levels and moisture stays roughly constant but with greater variability with increasing depth.

Contrasts between the dates for each treatment's final slope give 95%credible intervals which include zero for all treatments, except for treatments 1 and 9, which each show 4 of the 10 possible differences between the dates' final slopes as differing.

If we group long fallowing, response cropping and pastures and calculate a common final slope for each grouping, these estimated slopes all have 95% credible intervals which include zero. Comparing the contrasts for these grouped slopes across dates, no differences are found across dates. Mean moisture levels from 200 cm to 300 cm do not change across the different dates for the various treatments and types of cropping, but become more variable with depth.

4.5 *Point by point contrasts*

Figures 4- 6 graph the point by point contrasts for all depths and dates. The most important of these, Figure 4, shows the 95% credible intervals for the long fallow versus response cropping contrast as generally differing across dates at the shallower depths, but overlapping for the depths from 200 cm to 300 cm.

Tables for contrasts for the three-knot linear model are not shown but are discussed below. Table 5 shows the sign of each contrast whose 95% credible interval does not contain zero.

The statistical evidence is that the treatments no longer affect the moisture values from the depth of 200 cm to 300 cm, and given that the moisture profile is effectively flat at these depths, it seems that moisture levels after 200 cm are constant for their treatment, but have greater variability than at the mid-depths. The contrast of long fallow cropping (treatments 1-3) versus response cropping (treatments 5 & 6), has almost positive 95% credible intervals from 200cm to 300cm for all dates. Thus, it would appear that for the five dates considered response cropping decreases moisture levels at the depth critical for salination.

As expected, all contrasts for the contrast ‘Crop vs pasture’ (the average of treatments 1-6 minus the average of treatments 7-9) are positive for all dates and depths, with the difference being roughly constant from 200cm to 300 cm. That is, cropping leads to moister soil than pastures.

The lucerne pasture mixtures (treatments 7 & 8) perform consistently better than the native pastures for depths greater than 100 cm. That is, at these depths, lucerne mixtures lead to drier soil than the native pastures.

The differences discussed above are also shown in the saturated model contrast differences but not so markedly. These same contrasts when compared across the dates show essentially no differences in the depths from 200 cm to 300 cm.

4.6 *Spatial residual components, ψ*

As indicated in Section 3.4, no formal comparisons were made for the spatial residuals across dates. Graphs of their 95% credible intervals were plotted to inspect overlap or non-overlap. For depths from 140 cm and deeper the credible intervals largely overlap. Figure 7 gives contour graphs for these spatial residuals at the

depth of 240 cm for the different dates. These show considerable consistency across dates.

5. Discussion

In considering longitudinal agricultural experiments, Piepho et al. [27], Piepho and Ogutu [28], Piepho et al. [29], Wang and Goonewardene [41] and Brien and Demetrio [12] use mixed models within a REML framework to analyse their spatio-temporal data, and explicitly address the fitting of state-space models via standard software and REML. The fixed part of their models is generally simple and the data are measured on two spatial dimensions. Some soil profile studies [22] do not use spatial information in the analysis. Some studies composite the soils from different depths across soil types or treatment [38]. Others [24] use the mixed modelling framework advocated by Piepho et al. [27]. Roy and Blois [34] is one of the few papers in an agricultural context which uses conditional autoregressive models. For agricultural data, the current methodology of choice to account for spatial correlation would seem to be mixed modelling to describe spatial and other variance components, using REML. Despite the work of Besag et al. [5], Besag and Higdon [6, 7] there has been almost no use of CAR models for agricultural analyses. We use conditional autoregressive models for their simplicity and their capacity to allow reasonably complex fixed model components. Working with the sparse precision matrix from the adjacency matrix rather than from a dense covariance matrix permits efficient model fitting. Besag and Mondal [8], Lindgren et al. [20] show the equivalence of various kriging and CAR models.

The use of block updating allows good mixing and the Krylov subspace methods exploit the sparse structure of the precision matrix to give efficient sampling.

The choice to allow neighbours only at the same depth (the layered CAR model) is made for several reasons. Firstly, with depth an important part of the regression component, to include depth-neighbours would confound estimation of the treatment effects. Secondly, and more importantly, it permits the fitting of differing variances for the spatial components at each depth. Finally, using the obvious choice of distance weighted neighbours would mean that with the great differences in scale between horizontal and vertical distances the neighbourhood model would

degenerate effectively into a depth neighbourhood model only, while using (1,0) neighbours would also be difficult to justify. This consideration seems likely to apply in many agricultural contexts where observations are made in three spatial dimensions. We use a first order neighbourhood across the horizontal lattice and (1,0) weights.

We fitted the same model to five dates of data aiming to discover how best to fit a model for the full data. It largely appears that several important parameters of the model (ρ , τ and σ) are constant across dates for the depths which are of concern for salination.

From the DIC values, we see that the simplification of the three-knot model, where a longer linear segment at the deeper depths is used, has resulted in a better model. Clearly for depths from about 200 cm and greater, the various treatments no longer exercise a direct effect on the moisture content of the soil. Rather, the moisture content remains approximately constant at whatever level it has reached by 200 cm, but with increasing variability with increasing depth. This is true for all five dates.

We have presented a methodology for the analysis of three dimensional lattice data sets, where the distance between lattice points in one dimension is not commensurate with those in the other two, a situation which often applies water column, air column and soil studies. We call this model the layered CAR model. The layering is applicable to both regular and irregular lattices in the horizontal plane. We see layered models as applying to oceanographic, and air column data as well as three dimensional agricultural studies.

The analyses of the case study here have uncovered important features of the data. In particular, by having taken out the spatially correlated components, they indicate that response cropping gives rise to more satisfactory moisture levels than long fallow cropping below the root zone where the soils are at greatest risk of salination.

6. Tables

Table 1. Summary of DICs

	<i>Model</i>	<i>pD</i>	<i>DIC</i>	
Date 1	D135(S)	1065	-5850	
	K5(S)	1019	-5875	
	K3(S)	1032	-5915	*
	K3(NS)	959	-5657	
Date 2	D135(S)	1061	-5952	
	K5(S)	1039	-6038	
	K3(S)	1048	-6049	*
	K3(NS)	904	-5560	
Date 3	D135(S)	1049	-5885	
	K5(S)	1034	-5961	
	K3(S)	1044	-5996	*
	K3(NS)	906	-5509	
Date 4	D135(S)	1093	-6570	
	K5(S)	1070	-6623	*
	K3(S)	1064	-6619	
	K3(NS)	1011	-6507	
Date 5	D135(S)	1053	-6321	
	K5(S)	1024	-6378	
	K3(S)	1024	-6396	*
	K3(NS)	973	-6214	
<i>D135(S): Saturated model, 9×15 fixed terms.</i>				
<i>K5(S): 5-knot linear spline, 9×7 fixed terms.</i>				
<i>K3(S): 3-knot linear spline, 9×5 fixed terms.</i>				
<i>K3(NS): 3-knot linear spline - Intrinsic CAR.</i>				
<i>(S): Stationary CAR</i>				
<i>(NS): Non-stationary (Intrinsic CAR).</i>				
<i>pD: Estimated number of parameters.</i>				

Table 2. Estimates for $\hat{\rho}$ in the spatial precision matrix

	$\hat{\rho}$	95% CI
July 23, 1997	.461	(.373, .550)
December 4, 1997	.385	(.304, .477)
April 28, 1998	.375	(.289, .462)
September 23, 1998	.346	(.266, .429)
February 25, 1999	.325	(.246, .413)

Table 3. Differences in $\hat{\rho}$ across the five dates.

Date1	Date2	Est	q025	q975	Sig
1	2	0.077	0.000	0.156	*
	3	0.087	0.008	0.162	*
	4	0.115	0.032	0.195	*
	5	0.136	0.058	0.213	*
2	3	0.010	-0.068	0.085	
	4	0.039	-0.039	0.122	
	5	0.060	-0.017	0.135	
3	4	0.029	-0.049	0.114	
	5	0.050	-0.024	0.124	
4	5	0.021	-0.054	0.098	
$Est = \hat{\rho}_{Date1} - \hat{\rho}_{Date2}$					

Table 4. Slopes for segment 200 cm - 300 cm for each treatment

<i>Treatment</i>	<i>Day (Date)</i>	<i>Est</i>	<i>q025</i>	<i>q975</i>	<i>Sig</i>
1	1	-0.002	-0.010	0.005	
	2	0.001	-0.006	0.008	
	3	-0.001	-0.008	0.005	
	4	-0.005	-0.012	0.001	
	5	-0.002	-0.009	0.005	
2	1	-0.005	-0.013	0.002	
	2	-0.004	-0.011	0.003	
	3	0.010	0.003	0.017	*
	4	0.012	0.005	0.019	*
	5	-0.003	-0.010	0.003	
3	1	0.006	-0.001	0.014	
	2	0.004	-0.003	0.011	
	3	0.005	-0.002	0.012	
	4	0.002	-0.005	0.009	
	5	0.001	-0.006	0.008	
4	1	-0.005	-0.012	0.002	
	2	-0.004	-0.010	0.003	
	3	-0.004	-0.011	0.002	
	4	-0.004	-0.010	0.002	
	5	-0.005	-0.011	0.001	
5	1	0.001	-0.007	0.008	
	2	-0.000	-0.007	0.007	
	3	0.002	-0.005	0.009	
	4	0.000	-0.006	0.007	
	5	-0.000	-0.007	0.006	
6	1	0.004	-0.004	0.012	
	2	0.001	-0.006	0.008	
	3	0.005	-0.003	0.012	
	4	0.003	-0.004	0.010	
	5	-0.004	-0.011	0.003	
7	1	0.003	-0.003	0.010	
	2	0.003	-0.003	0.010	
	3	0.005	-0.002	0.011	
	4	0.003	-0.003	0.009	
	5	0.002	-0.004	0.008	
8	1	-0.009	-0.016	-0.003	*
	2	-0.009	-0.016	-0.002	*
	3	-0.009	-0.015	-0.003	*
	4	-0.008	-0.014	-0.002	*
	5	-0.011	-0.017	-0.005	*
9	1	0.004	-0.003	0.011	
	2	0.003	-0.004	0.010	
	3	0.006	-0.001	0.013	
	4	0.006	-0.001	0.012	
	5	0.005	-0.001	0.012	

* indicates 95% credible interval does not include zero.

Table 5. Signs for contrasts with 95% credible intervals not including zero, for each date. Positive (+) and negative (−) values indicated.

<i>Contrast</i>	<i>Depth</i>	<i>Date</i>				
		<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>
Long Fallow - Response	20	+	-	+		+
	40	+		+	+	+
	60	+		+		+
	80	+	+	+		+
	100	+	+	+	+	+
	120	+	+		+	+
	140	+		-	+	+
	160	+	+		+	+
	180	+	+		+	+
	200	+	+		+	+
	220	+	+	+	+	+
	240	+	+	+	+	+
	260	+	+	+	+	+
	280				+	+
	300					+
Cropping - Pastures	20	+	+	+	+	+
	40	+	+	+	+	+
	60	+	+	+	+	+
	80	+	+	+	+	+
	100	+	+	+	+	+
	120	+	+	+	+	+
	140	+	+	+	+	+
	160	+	+	+	+	+
	180	+	+	+	+	+
	200	+	+	+	+	+
	220	+	+	+	+	+
	240	+	+	+	+	+
	260	+	+	+	+	+
	280	+	+	+	+	+
	300	+	+	+	+	+
Lucerne mixtures - Native	20				-	-
	40	-			-	-
	60	-	-		-	-
	80	-	-		-	-
	100	-	-	-	-	-
	120	-	-	-	-	-
	140	-	-	-	-	-
	160	-	-	-	-	-
	180	-	-	-	-	-
	200	-	-	-	-	-
	220	-	-	-	-	-
	240	-	-	-	-	-
	260	-	-	-	-	-
	280	-	-	-	-	-
	300	-	-	-	-	-

Figure 1. Relative Deviance curves for date 4, September 23, 1998. The solid line represents the deviance for the saturated model, the middle broken line the deviance for the 3-knot linear spline, and the more coarsely broken line on the left the deviance for the 5-knot linear spline model.

7. Figures

Figure 2. Square root of non-spatial variance components, by date and depth. Credible intervals are staggered in date order.

Figure 3. Square root of spatial variance components, by date and depth. Credible intervals are staggered in date order.

Figure 4. Contrast: Long Fallow - Response cropping. Credible intervals are staggered in date order.

Figure 5. Contrast: Cropping - Pastures. Credible intervals are staggered in date order.

Figure 6. Contrast: Lucerne mixtures - Native pastures. Credible intervals are staggered in date order.

Figure 7. Spatial residual components at depth 240 cm.

References

- [1] Aitkin, M.: 1997, The calibration of P-values, posterior Bayes factors and the AIC from the posterior distribution of the likelihood, *Statistics and Computing* **7**, 253–261.
- [2] Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Croz, J. D., Greenbaum, A., Hammarling, S., McKenney, A. and Sorensen, D.: 1999, *LAPACK Users' Guide: Third Edition*, 22 aug 1999 edn, Society for Industrial and Applied Mathematics (SIAM), Philadelphia.
URL: <http://www.netlib.org/lapack/>
- [3] Belitz, C., Brezger, A., Kneib, T. and Lang, S.: 2009a, BayesX Software for Bayesian Inference in Structured Additive Regression Models Version 2.0.1 Reference Manual, Online at <http://www.stat.uni-muenchen.de/bayesx/bayesx.html>. Accessed: October 25, 2010.
- [4] Belitz, C., Brezger, A., Kneib, T. and Lang, S.: 2009b, BayesX Software for Bayesian Inference in Structured Additive Regression Models Version 2.0.1 Software Methodology Manual, Online at <http://www.stat.uni-muenchen.de/bayesx/bayesx.html>. Accessed: October 25, 2010.
- [5] Besag, J. E., Green, P., Higdon, D. and Mengersen, K.: 1995, Bayesian computation and stochastic systems, *Statistical Science* **10**(1), 3–41.
- [6] Besag, J. E. and Higdon, D.: 1993, Bayesian inference for agricultural field experiments, *Bull. Inst. Internat. Statist* **55**(Book 1), 121–136.
URL: <http://www.stat.duke.edu/higdon/Papers/isi1.ps>
- [7] Besag, J. E. and Higdon, D.: 1999, Bayesian analysis of agricultural field experiments, *Journal of the Royal Statistical Society Series B-Statistical Methodology* **61**, 691–717. Part 4.
- [8] Besag, J. E. and Mondal, D.: 2005, First-order intrinsic autoregressions and the de Wijs process, *Biometrika* **92**(4), 909–920.
- [9] Besag, J. E., York, J. and Mollié, A.: 1991, Bayesian image restoration with applications in spatial statistics (with discussion), *Annals of the Institute of Mathematical Statistics* **43**, 1–59.
- [10] Blackford, L., Demmel, J., Dongarra, J., Duff, I., Hammarling, S., Henry, G., Heroux, M., Kaufman, L., Lumsdaine, A. and Petitet, A.: 2002, An updated set of basic linear algebra subprograms (BLAS), *ACM Transactions on Mathematical Software (TOMS)* **28**(2), 135–151.
URL: <http://www.netlib.org/blas/>
- [11] Brezger, A. and Lang, S.: 2006, Generalized structured additive regression based on Bayesian P-splines, *Computational Statistics and Data Analysis* **50**(4), 967–991.
- [12] Brien, C. J. and Demetrio, C. G. B.: 2009, Formulating mixed models for experiments, including longitudinal experiments, *Journal of Agricultural, Biological, and Environmental Statistics* **14**(3), 253–280.
- [13] Chib, S. and Carlin, B. P.: 1999, On MCMC sampling in hierarchical longitudinal models, *Statistics and Computing* **9**, 17–26.
- [14] Donald, M., Alston, C. L., Young, R. R. and Mengersen, K. L.: 2011, A bayesian analysis of an agricultural field trial with three spatial dimensions, *Computational Statistics and Data Analysis* **55**(12), 3320–3332. Available online 25 June 2011.
- [15] Gelfand, A. E. and Vounatsou, P.: 2003, Proper multivariate conditional autoregressive models for spatial data analysis, *Biostatistics* **4**(1), 11–25.
- [16] Geweke, J.: 1992, Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, *Bayesian Statistics* **4**, 169–188.
- [17] Hastie, T. and Tibshirani, R.: 1990, *Generalized additive models*, Monographs on statistics and applied probability, 1st edn, Chapman and Hall, London ; New York.
- [18] Kneib, T. and Fahrmeir, L.: 2006, Structured additive regression for categorical spacetime data: A mixed model approach, *Biometrics* **62**(1), 109–118.

- [19] Lang, S. and Brezger, A.: 2004, Bayesian P-splines, *Journal of Computational and Graphical Statistics* **13**(1), 183–212.
- [20] Lindgren, F., Rue, H. and Lindstrom, J.: 2011, An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.
- [21] Lui, J. S., Wong, W. H. and Kong, A.: 1994, Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentations schemes., *Journal of the Royal Statistical Society, Series B* **57**(1), 157–169.
- [22] Macdonald, B. C. T., Reynolds, J. K., Kinsela, A. S., Reilly, R. J., van Oploo, P., Waite, T. D. and White, I.: 2009, Critical coagulation in sulfidic sediments from an east-coast Australian acid sulfate landscape, *Applied Clay Science* **46**(2), 166–175.
- [23] Martino, S. and Rue, H.: 2009, *R Package: INLA*, Department of Mathematical Sciences NTNU, Norway.
- [24] Nayyar, A., Hamel, C., Lafond, G., Gossen, B. D., Hanson, K. and Germida, J.: 2009, Soil microbial quality associated with yield reduction in continuous-pea, *Applied Soil Ecology* **43**(1), 115–121.
- [25] Ngo, L. and Wand, M.: 2004, Smoothing with mixed model software, *Journal of Statistical Software* **9**, 1–56.
- [26] NumPy Community: 2010, NumPy Reference Manual: Release 1.5.0.dev8106, Available online: <http://docs.scipy.org/doc/>. Accessed: February 9, 2010.
URL: <http://docs.scipy.org/doc/>
- [27] Piepho, H. P., Buchse, A. and Richter, C.: 2004, A mixed modelling approach for randomized experiments with repeated measures, *Journal of Agronomy and Crop Science* **190**(4), 230–247.
- [28] Piepho, H. P. and Ogutu, J. O.: 2007, Simple state-space models in a mixed model framework, *American Statistician* **61**(3), 224–232.
- [29] Piepho, H. P., Richter, C. and Williams, E.: 2008, Nearest neighbour adjustment and linear variance models in plant breeding trials, *Biometrical Journal* **50**(2), 164–189.
- [30] Pitt, M. and Shephard, N.: 1999, Analytic convergence rates and parameterisation issues for the Gibbs sampler applied to state space models., *Journal of Time Series Analysis* **20**, 63–85.
- [31] Plummer, M., Best, N., Cowles, K. and Vines, K.: 2010, Coda R Package: Output analysis and diagnostics for MCMC.
- [32] Raftery, A. and Lewis, S.: 1992, How many iterations in the Gibbs sampler?, in J. Bernardo, J. Berger, A. Dawid and A. Smith (eds), *Bayesian Statistics 4*, Oxford University Press, Oxford.
- [33] Ringrose-Voase, A., Young, R. R., Payder, Z., Huth, N., Bernardi, A., Cresswell, H., Keating, B., Scott, J., Stauffacher, M., Banks, R., Holland, J., Johnston, R., Green, T., Gregory, L., Daniells, I., Farquharson, R., Drinkwater, R., Heidenreich, S. and Donaldson, S.: 2003, Deep drainage under different land uses in the Liverpool Plains Catchment, *Technical Report 3*, Agricultural Resource Management Report Series, NSW Agriculture Orange.
- [34] Roy, V. and Blois, S. d.: 2008, Evaluating hedgerow corridors for the conservation of native forest herb diversity, *Biological Conservation* **141**, 298–307.
- [35] Rue, H. and Held, L.: 2005, *Gaussian Markov random fields : theory and applications*, Chapman & Hall/CRC, Boca Raton.
- [36] Saad, Y.: 2003, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics. [electronic resource].
URL: <http://gateway.library.qut.edu.au/login?url=http://www.igpublish.com/siam-ebook/search.nsp?query0=0-89871-534-2&field0=ISBN>
- [37] Simpson, D. P., Turner, I. W. and Pettitt, A. N.: 2008, Fast sampling form a Gaussian Markov random field using Krylov subspace approaches, *QUT Eprints 14376* (Brisbane), 1–17. Available

- online: <http://eprints.qut.edu.au>.
- [38] Sleutel, S., Vandenbruwane, J., De Schrijver, A., Wuyts, K., Moeskops, B., Verheyen, K. and De Neve, S.: 2009, Patterns of dissolved organic carbon and nitrogen fluxes in deciduous and coniferous forests under historic high nitrogen deposition, *Biogeosciences* **6**(12), 2743–2758.
 - [39] Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and van der Linde, A.: 2002, Bayesian measures of model complexity and fit, *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **64**(4), 583–639.
 - [40] Strickland, C. M., Simpson, D. P., Turner, I. W., Denham, R. and Mengersen, K. L.: 2011, Fast Bayesian analysis of spatial dynamic factor models for multitemporal remotely sensed imagery, *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **60**(1), 109–124.
 - [41] Wang, L. A. and Goonewardene, Z.: 2004, The use of mixed models in the analysis of animal experiments with repeated measures data, *Canadian Journal of Animal Science* **84**(1), 1–11.

8. Appendix: A Gibbs sampler for the CAR layered model

The model and its parameters are described in Section 3.1. The model comprises a regression component with a design matrix, X , and is a layered CAR model, in that the residuals (both spatial and random) are drawn from variance components specific to each depth. Additionally, the neighbourhood matrix permits neighbours only in the common depth layer, giving rise to a block diagonal neighbourhood matrix for the full data. For definitions of the parameters, we refer the reader to Section 3.1. Here we describe the Gibbs sampler used for their estimation.

The joint posterior for the full set of unknown parameters is estimated by partitioning the parameters into five blocks

$$(\psi, \tau, \beta, \sigma, \rho).$$

and a Gibbs sampling scheme is defined such that the j^{th} step is

- (1) Sample ψ^j from $p(\psi|\mathbf{y}, \tau^{j-1}, \beta^{j-1}, \sigma^{j-1}, \rho^{j-1})$,
- (2) Sample τ^j from $p(\tau|\mathbf{y}, \psi^j, \beta^{j-1}, \sigma^{j-1}, \rho^{j-1})$,
- (3) Sample β^j from $p(\beta|\mathbf{y}, \psi^j, \tau^j, \sigma^{j-1}, \rho^{j-1})$,
- (4) Sample σ^j from $p(\sigma|\mathbf{y}, \psi^j, \tau^j, \beta^j, \rho^{j-1})$,
- (5) Sample ρ^j from $p(\rho|\mathbf{y}, \psi^j, \tau^j, \beta^j, \sigma^j)$.

Let S be the number of horizontal sites, D the number of different depths, and n the number of observations.

The following subsections describe the sampling from each of the full conditional posteriors in the scheme above.

8.1 Sampling β .

We define $\tilde{\mathbf{y}}$, such that

$$\tilde{\mathbf{y}} = \Sigma^{-1/2}(\mathbf{y} - \psi),$$

and $\mathbf{X}$ such that

$$\tilde{\mathbf{X}} = \Sigma^{-1/2}\mathbf{X}.$$

The prior probability density function (pdf) for β is taken as

$$\beta \sim N(\underline{\beta}, \underline{V}^{-1}),$$

where $\underline{\beta}$ is the prior mean, and $\underline{V}$ is the prior precision. Thus the posterior distribution for β is given by

$$\beta | y, \sigma, \psi \sim N(\bar{\beta}, \bar{V}^{-1}),$$

where

$$\bar{V} = \tilde{X}^T \tilde{X} + \underline{V},$$

and

$$\bar{V}\bar{\beta} = \underline{V}\underline{\beta} + \tilde{X}^T \tilde{y}.$$

8.2 Sampling σ .

Let

$$\epsilon = y - X\beta - \psi,$$

and let the $n \times 1$ residual vector be partitioned by depth, d , into D subvectors, ϵ_d , such that

$$\epsilon = [\epsilon_1, \epsilon_2, \dots, \epsilon_D],$$

where each ϵ_d is an $S \times 1$ vector, with $d = 1, 2, \dots, D$. The vector σ^2 is updated by updating each variance component, σ_d^2 , one at a time, using the following updating equations.

$$1/\sigma_d^2 | y, \beta, \psi \sim \text{Gamma}(\bar{\nu}/2, \bar{s}/2),$$

where

$$\bar{\nu} = S + \underline{\nu},$$

and

$$\bar{s} = \underline{s} + \boldsymbol{\epsilon}_d^T \boldsymbol{\epsilon}_d,$$

and the common prior for each variance component, σ_d^2 , the d th element of the vector, $\boldsymbol{\sigma}^2$, is

$$1/\sigma_d^2 \sim \text{Gamma}(\underline{\nu}/2, \underline{s}/2), \quad d = 1, 2, \dots, D.$$

8.3 Sampling ψ .

Define

$$\tilde{\mathbf{y}} = \mathbf{y} - \mathbf{X}\boldsymbol{\beta}.$$

This gives

$$\tilde{\mathbf{y}}|\boldsymbol{\beta}, \boldsymbol{\sigma}, \rho, \boldsymbol{\tau} \sim N(\boldsymbol{\psi}, \boldsymbol{\Sigma}).$$

Hence,

$$\begin{aligned} p(\boldsymbol{\psi}|\mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\sigma}, \rho, \boldsymbol{\tau}) &\propto p(\tilde{\mathbf{y}}|\dots) \times p(\boldsymbol{\psi}|\rho, \boldsymbol{\tau}), \\ &\propto \exp \left\{ -\frac{1}{2} \left(\boldsymbol{\psi}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\psi} + \boldsymbol{\psi}^T \boldsymbol{\Omega} \boldsymbol{\psi} - 2\boldsymbol{\psi}^T \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{y}} \right) \right\}. \end{aligned}$$

and thus

$$\boldsymbol{\psi}|\mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\sigma} \sim N(\bar{\boldsymbol{\psi}}, \bar{\boldsymbol{\Omega}}^{-1}),$$

where

$$\bar{\boldsymbol{\Omega}} = \boldsymbol{\Omega} + \boldsymbol{\Sigma}^{-1},$$

and

$$\bar{\boldsymbol{\Omega}}\bar{\boldsymbol{\psi}} = \boldsymbol{\Sigma}^{-1}\tilde{\mathbf{y}}.$$

8.4 Sampling $\boldsymbol{\tau}$.

The elements τ_d^2 of the vector $\boldsymbol{\tau}^2$, $d = 1, 2, \dots, D$, are updated one at a time as follows. The $n \times 1$ vector $\boldsymbol{\psi}$ is partitioned into D subvectors $\boldsymbol{\psi}_d$, (the spatial

residuals at depth d), and (see Section 3.1)

$$\mathbf{\Omega}(\rho, \boldsymbol{\tau}) = \text{block diagonal}(\mathbf{Q}/\tau_1^2, \mathbf{Q}/\tau_2^2, \dots, \mathbf{Q}/\tau_D^2).$$

Let the prior pdf for τ_d^2 be given by

$$1/\tau_d^2 \sim \text{Gamma}\left(\frac{a}{2}, \frac{b}{2}\right),$$

with $(\underline{a}, \underline{b})$ as hyperpriors.

This gives the updating posterior probability density function for $\boldsymbol{\tau}^2$ as

$$1/\tau_d^2 | \boldsymbol{\psi}, \rho \sim \text{Gamma}\left(\frac{\bar{a}}{2}, \frac{\bar{b}}{2}\right), \text{ where}$$

$$\bar{a} = \underline{a} + S, \text{ and}$$

$$\bar{b} = \underline{b} + \boldsymbol{\psi}_d^T \mathbf{Q} \boldsymbol{\psi}_d,$$

for

$$\boldsymbol{\tau}^2 = [\tau_1^2, \tau_2^2, \dots, \tau_D^2].$$

8.5 Sampling ρ .

From Section 3.1, the precision matrix for the spatial components is $\mathbf{\Omega}(\rho, \boldsymbol{\tau})$. The prior for ρ is taken as

$$\rho \sim \text{Beta}(\alpha, \beta),$$

with (α, β) the hyperparameters.

Hence, the posterior probability density function for ρ is given by

$$p(\rho | \boldsymbol{\psi}, \boldsymbol{\tau}) \propto |\mathbf{\Omega}(\rho, \boldsymbol{\tau})|^{1/2} \rho^{\alpha-1} (\rho-1)^{\beta-1} \exp\left\{-\frac{1}{2}(\boldsymbol{\psi}^T \mathbf{\Omega}(\rho, \boldsymbol{\tau}) \boldsymbol{\psi})\right\},$$

and ρ is sampled via a Metropolis-Hastings update.