

The use of systems of stochastic PDEs as priors for multivariate models with discrete structures

Erlend Aune and Daniel P. Simpson
Department of Mathematical Sciences, NTNU, Norway

August 9, 2012

Abstract

A challenge in multivariate problems with discrete structures is the inclusion of prior information that may differ in each separate structure. A particular example of this is seismic amplitude versus angle (AVA) inversion to elastic parameters, where the discrete structures are geologic layers. Recently, the use of systems of linear stochastic partial differential equations (SPDEs) have become a popular tool for specifying priors in latent Gaussian models. This approach allows for flexible incorporation of nonstationarity and anisotropy in the prior model. Another advantage is that the prior field is Markovian and therefore the precision matrix is very sparse, introducing huge computational and memory benefits. We present a novel approach for parametrising correlations that differ in the different discrete structures, and additionally a geodesic blending approach for quantifying fuzziness of interfaces between the structures.

Keywords: Gaussian distribution, multivariate, stochastic PDEs, discrete structures

1 Introduction

In spatial statistics, the need for specifying different behaviour in different regions in space is crucial for making a good prior model. The literature is abundant with methodologies for this. In the multivariate setting, this generalises to having different correlations between the fields in different regions and different cross-differentiability properties.

A particular model problem where this is important is the seismic AVA inversion problem, which well studied in the geophysical literature. There are several incarnations of this problem with varying degrees of complexity. In this article, our primary example is the inversion problem studied in Buland and Omre (2003); Buland et al. (2003); Rabben et al. (2008), using the wavefield propagation approximations in Aki and Richards (1980), which results in linear systems of equations to solve. Variants and extensions of these equations are found in Stovas and Ursin (2003), including nonlinear approximations that may yield better inversion results in some situations. We exemplify our contributions using this example explicitly throughout the article.

The model we adopt in this text is the same as in Buland and Omre (2003), which is essentially

$$d(\mathbf{s}) = w \star r_{PP}(\mathbf{m})(\mathbf{s}) + \epsilon, \quad (1)$$

where \star denotes convolution in time, w is an approximation to the source wavelet – i.e. the shape of the wave traveling through the subsurface, and $r_{PP}(\mathbf{m})$ denotes a reflectivity operator. The reflectivity operator takes relative differences in elastic parameters to reflection

coefficients for the wave. We adopt the following elastic parameters,

$$m_1 = \frac{\Delta v_P}{\bar{v}_P}, \quad m_2 = \frac{\Delta v_S}{\bar{v}_S}, \quad m_3 = \frac{\Delta \rho}{\bar{\rho}}. \quad (2)$$

I.e. m_1, m_2, m_3 denotes the relative difference of P -velocity (pressure wave velocity), S -velocity (shear wave velocity) and density respectively, and the reflectivity operator is defined by

$$r_{PP}(\theta) = \frac{m_1}{2}(1 + \tan^2 \theta) - 4m_2\gamma^2 \sin^2 \theta + \frac{m_3}{2}(1 - 4\gamma^2 \sin^2 \theta), \quad (3)$$

there θ is the reflection angle and γ^2 denotes a background (v_S/v_P)-ratio. Rewriting this in matrix notation yields

$$\mathbf{d} = \mathbf{W}\mathbf{A}\mathbf{m} + \boldsymbol{\epsilon}, \quad (4)$$

where \mathbf{d} are observations, \mathbf{W} is the discretized wavelet operator, \mathbf{A} , the discretized reflectivity operator, \mathbf{m} the elastic parameters and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ an error term which is often assumed to be normally distributed.

In this text, we will explore a novel method for designing a good prior for \mathbf{m} using linear systems of stochastic partial differential equations. We emphasize, however, that while the approach developed here is designed with seismic AVA inversion in mind, it is very flexible and can be adopted in any setting where we have multivariate fields with separate regions where we would like to incorporate prior information.

All the figures that appear in this text have comparative scales, so that the colour schemes have the same min-max values in each individual figure. Hence, the figures makes sense, without cluttering them with additional colourbars.

2 Prior specification

The choice of prior in the inversion problem is of great importance when it comes to the performance of the inversion. It is vital to choose a “good” prior to emphasise the properties of \mathbf{m} that we know it has. For us, \mathbf{m} will denote the parameters of interest, and it depends on position. We construct the prior by combining heuristics and expert knowledge of the spatial model. For a Gaussian prior model, the standard way of specifying the prior model is through the covariance function, which is often assumed to be stationary (see, e.g. Buland and Omre (2003)). A stationary covariance function is defined by a correlation function that defines how much a point is correlated with its neighbours and a marginal variance parameter, ϱ^2 through

$$\varrho^2 c(\|\mathbf{x} - \mathbf{y}\|_{\mathbf{A}}) = \text{Cov}(\mathbf{x}, \mathbf{y}), \quad (5)$$

where \mathbf{A} is a positive definite matrix that defines the non-changing anisotropy of the field. In the Gaussian case, this defines a strictly stationary process if the mean is constant. There is a list of widely used covariance functions in Cressie (1993). We will throughout this text assume that the prior is from the Gaussian family. This family is defined by having density

$$p(\mathbf{x}|\mathbf{Q}, \boldsymbol{\mu}_x) = (2\pi)^{-n/2} \det(\mathbf{Q})^{1/2} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_x)^T \mathbf{Q}(\mathbf{x} - \boldsymbol{\mu}_x)\right), \quad (6)$$

where $\mathbf{Q} = \boldsymbol{\Sigma}^{-1}$ is the precision matrix – the inverse of the covariance matrix $\boldsymbol{\Sigma}$ – and $\boldsymbol{\mu}_x$ is the expectation, $\mathbb{E}(\mathbf{x}|\boldsymbol{\mu}_x)$.

Moreover, the fields m_1, m_2, m_3 are assumed correlated with correlations specified by well data and/or other local knowledge. In the discretized domain, this allows for the following decomposition of the total covariance matrix

$$\boldsymbol{\Sigma}_m = \boldsymbol{\Sigma}_{\text{space}} \otimes \boldsymbol{\Sigma}_0 \quad (7)$$

where Σ_{space} denotes the spatial covariance matrix, typically defined through a covariance function, and Σ_0 the correlations between the elastic parameters. Since seismic observations typically are on a regular grid, either in 2-D or 3-D, it is possible to let Σ_{space} be circulant by extending the grid by as many points as is needed to get the correlation below a threshold – typically 0.1 or 0.05. This allows us to use fast Fourier transforms for computing quantities of interest related to the covariance matrix. This, together with the Kronecker structure of Σ_m allows for fast computations. See Buland et al. (2003); Rue and Held (2005); Gray (2006) for details. This approach also has very low memory requirements; since Σ_{space} is circulant it may be stored using only one vector. Hence storage is $\mathcal{O}(n)$ and computations (of any kind) are at most $\mathcal{O}(n \log n)$, where n is the number of nodes in the extended lattice.

2.1 SPDE formulation

While this decomposition is sensible, it is also very inflexible and requires stationarity for low storage requirements. Another way of pursuing good prior models with fast computations and low memory requirements is through the use of elliptic (pseudo) differential operators (Ruzhansky and Turunen (2009), part 2 is an accessible source). The theory of pseudo differential operators is closely related to Weyl transforms and short-time Fourier transforms or Gabor transforms (Feichtinger et al. (2008)) and usual spectral considerations in seismology apply. In this approach, it is the sparsity of the resulting precision matrices that makes storage and computation manageable. Recently, Lindgren et al. (2011), studied how to apply such operators in a statistical setting. They studied a Riesz-Bessel operator, $(-\Delta + \kappa)^{\alpha/2}$ and its relation to computation and Matérn covariance models (Matérn, 1960; Whittle, 1963). The main lessons are firstly, if

$$M_{\kappa, \alpha} x(s) := (\kappa^2 - \Delta)^{\alpha/2} x(s) = \mathcal{W}(s), \quad (8)$$

where \mathcal{W} is spatial Gaussian white noise, then $x(s)$ has Matérn type covariance function, i.e.,

$$\rho(r) = \frac{\varrho^2}{\Gamma(\alpha - d/2) 2^{\alpha - d/2 - 1}} (\kappa r)^{\alpha - d/2} K_{\alpha - d/2}(\kappa r), \quad (9)$$

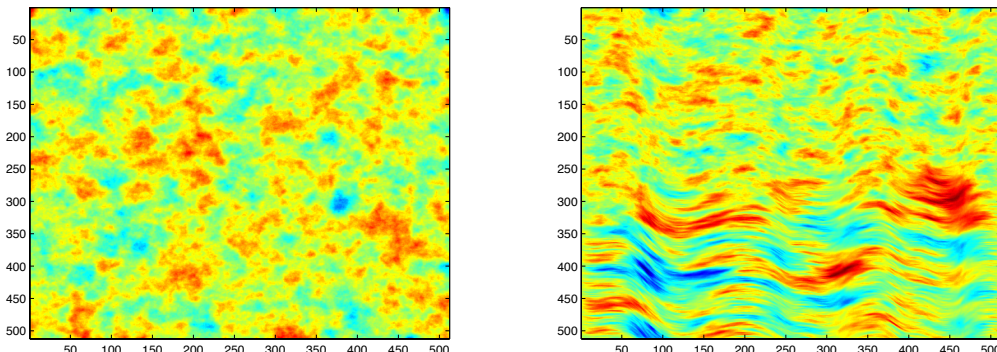
$$\varrho^2 = \frac{\Gamma(\alpha - d/2)}{\Gamma(\alpha) (4\pi)^{d/2} \kappa^{2(\alpha - d/2)}}, \quad (10)$$

where K_s is the modified Bessel function of the first kind. Secondly, fast computations through finite element methods or other discretisations of the differential operator in (8) are available through the induced Markov properties of the discretisation matrix, $\mathbf{Q}_{\text{space}}$. That essentially means that $\mathbf{Q}_m = \mathbf{Q}_{\text{space}} \otimes \mathbf{Q}_0$ is (very) sparse and with a structure amenable to Cholesky factorisation. An alternative requirement is that we can construct the matrix vector product $\mathbf{Q}_m \mathbf{v}$ and $\det(\mathbf{Q}_m)$ relatively quickly through some iterative or direct procedure, see Simpson (2008); Aune et al. (2012a,b)

When addressing the “stationarity” of the field defined by (8), it is only stationary in the sense of (5) if it is defined on the whole of \mathbb{R}^k , where $k = 2, 3$ in our case – alternatively when the corresponding operator is defined on a manifold without boundary. In our case the domain on which (5) is defined is merely a subset, namely a square or box in \mathbb{R}^2 or \mathbb{R}^3 . Hence boundary effects resulting from boundary conditions may destroy its direct interpretability in terms of this equation. It is, of course, possible to specify boundary conditions in such a way that you retain the property in (5), but usually there are more natural physical boundary conditions that in our opinion improves upon the specification through SPDEs compared to the model defined by covariance matrices through stationary covariance functions also in the stationary case.

There are two properties that are desirable to have in the prior model in AVA inversion. The first is being able to have different correlation length at different points in space. If a geologist has sound reasons to believe that a layer is very inhomogeneous, it may warrant putting a lower correlation length here than in a layer that is thought to be very homogeneous with very similar properties. Facilitating this is trivial – one merely lets $\kappa^2 = \kappa^2(s)$ vary with space. The other property that is very desirable to have is anisotropy. Letting the correlation

Figure 1: Realisations from stationary model given by (8) (left) and non-stationary model given by (11) (right). The non-stationary model has a curved interface, and the field below the interface has anisotropy directed along the curve, while above the interface it is almost isotropic.



length vary with direction is very natural given that the layers are typically not flat but are deformed in a specific way. The SPDE resulting is the following variant of (8):

$$M_{\kappa,\alpha,s}x(s) = (\kappa^2(s) - \nabla \cdot \mathbf{H}(s)\nabla)^{\alpha/2}x(s) = \mathcal{W}(s), \quad (11)$$

where \mathbf{A} is a 3×3 symmetric positive definite matrix defining the anisotropy angle and principal correlation length in the three directions defined by the eigenvectors of the matrix. Realisations of the stationary model and the nonstationary model is given in Figure 1. Here we have illustrated the “layer” flexibility mentioned above, where the top layer is isotropic, and the bottom layer is anisotropic with deformation defined by the layer.

To see how this relates to the usual approach, consider $\mathbf{Q}_0 = \Sigma_0^{-1}$ and say that m_1, m_2, m_3 have equal Matérn covariance models (this includes the widely used exponential and Gaussian models), then the prior given as in (7) is given by the following system of stochastic differential equations:

$$(M_{\kappa,\alpha,s} \otimes \mathbf{Q}_0)\mathbf{m} = \mathcal{W} \quad (12)$$

where \mathcal{W} is vector Gaussian spatial white noise. The experience in AVA-inversion is that at least one component of \mathbf{m} worse resolved than the others, with m_1 being resolved the best (see Rabben et al. (2008) or any other article treating this problem). The obvious next question then is whether or not (12) specifies the best way of lending strength to the least resolved parameters. If not, can we find better operators on the diagonal in (12), and/or replace the off-diagonals with other operators that have better properties in the inversion problem? The answer to this question is not obvious, but we investigate some alternatives and see how they perform in our inversion problem; the criterion for a better prior in the synthetic case being that $\mathbb{E}((\mathbf{m}_{\text{true}} - \mathbf{m}_{\text{est}}^{\text{new}})^2) < \mathbb{E}((\mathbf{m}_{\text{true}} - \mathbf{m}_{\text{est}}^{\text{base}})^2)$, where $\mathbf{m}_{\text{est}}^{\text{base}}$ is given by the prior model (7).

It is possible to replace the operator $M_{\kappa,\alpha}$ in (12) by more general pseudo-differential operators. Representations of such operators in terms of its symbol are given by

$$(K_{\sigma}f)(x) = \int_{\mathbb{R}^d} \sigma(x, \xi) \hat{f}(\xi) e^{2\pi i x \cdot \xi} d\xi, \quad (13)$$

where \hat{f} is the Fourier transform of f , and σ is the symbol of the operator. The symbol can be interpreted as defining the local spectrum of the operator. A deep theorem given in Rozanov (1977) states that a stationary random field is Markov (in the continuous sense) if and only if σ^{-1} is a symmetric positive polynomial. Hence Markov fields are represented by differential operators. Now, if the field in question is not Markov, it is possible to approximate σ by

a rational approximation, $\sigma(x, \xi)^{-1} \approx \sigma_{\text{rat}}^{-1}(x, \xi) = \sum_{j=0}^k a_j(x)(2\pi i \xi)^j$. To find the a_j s one can, for instance, use optimisation techniques. This is one way to do it, but we suspect that the time-frequency localisation of such an approach may be suboptimal, and discretization of the non-Markov operator may be better suited for time-frequency compressing approaches inducing approximate Markovity. We do not pursue these type of ideas here, but mention them as they are good candidates for future research.

3 Systems of SPDEs – generalising “ \mathbf{Q}_0 ”

It is easy to write the form the generalised approach must have. First, for $i, j = 1, \dots, 3$, let

$$K_{ij} = q_{ij}(\mathbf{s})(\kappa_{ij}(\mathbf{s}) - \nabla \cdot \mathbf{A}_{ij}(\mathbf{s})\nabla)^{\alpha_{ij}/2} \quad (14)$$

and define the following system of SPDEs

$$\mathbf{K}\mathbf{m}(\mathbf{s}) = \begin{pmatrix} K_{11} & K_{12} & K_{13} \\ K_{12} & K_{22} & K_{23} \\ K_{13} & K_{23} & K_{33} \end{pmatrix} \mathbf{m}(\mathbf{s}) = \mathcal{W}(\mathbf{s}) \quad (15)$$

For $q_{ij}(\mathbf{s}) = Q_{ij}^0$ and $K_{ij} = M_{\kappa, \alpha}$ we recover the structure in the previous section with stationarity. For convenience, we call $q_{ij}(\mathbf{s})$ the blending coefficients. In Hu et al. (2012), they study the properties of this model in the stationary case, and give the link to the multivariate Matérn fields in Gneiting et al. (2010). Any choice of K_{ij} defines a valid Gaussian Markov random field, both in the continuous sense and when discretized. In our treatment, we restrict ourselves to the case where $\alpha_{ij} = \alpha$, $\mathbf{A}_{ij}(\mathbf{s}) = \mathbf{A}(\mathbf{s})$ and $\kappa_{ij}(\mathbf{s}) = \kappa(\mathbf{s})$.

3.1 Parametrising the blending coefficients

In general, it is both hard to interpret a local precision matrix, $\mathbf{Q}_0(\mathbf{s}) = \{q_{ij}(\mathbf{s})\}_{ij}$ defining how the individual parts of the multivariate fields is related to each other at position \mathbf{s} , and to ensure that this matrix is positive definite. It is much more natural to work with the inverse, namely the correlation matrix defining the local correlation of the fields, $\mathbf{\Sigma}_0(\mathbf{s}) = \mathbf{Q}_0^{-1}(\mathbf{s})$. The $q_{ij}(\mathbf{s})$ is then simply given by the corresponding matrix elements. In the AVA inversion problem, information about correlation in different layers may come from geologists or geophysicists for who may know of phase changes when going from one layer to another in the different layers, or other, more complex phenomena. It may also come from well-logs that may contain information about such matters.

Suppose that $\mathbf{\Sigma}_0(\mathbf{s}) = \mathbf{\Sigma}_{0,1}$ for $\mathbf{s} \in S_1 \subset \mathbb{R}^d$ and $\mathbf{\Sigma}_{0,2}$ for $\mathbf{s} \in S_2 \subset \mathbb{R}^d$. Then we have a model that has specific correlations in one spatial region of the multivariate fields, and different correlations in another spatial region. There is obviously a transition between these two states. If the transition is discontinuous, this may be seen as a discontinuity of the correlations in the realisation of the multivariate random field, which may make sense in some situations.

In order to visualise what this means, we give realisations of the four major prior models we have discussed. In Figure 2, no prior information about the geometry of the subsurface can be included. In Figure 3, geometric information has been incorporated, but no change in the correlation between the parameters in space can be included. In Figure 4, an example realisation from the full model is given. Pay attention to the rightmost field – here the correlation to the other two fields changes from being positive in the top layer to being negative in the bottom layer.

3.2 Geodesic blending

There are obviously many ways of making a smooth transition between $\mathbf{\Sigma}_{0,1}$ and $\mathbf{\Sigma}_{0,2}$, but one key consideration is that $\mathbf{\Sigma}_0(\mathbf{s})$ must remain positive definite for all \mathbf{s} in some transition domain S_T . One thing is certain - it is not necessarily enough to let the off-diagonals element in $\mathbf{\Sigma}_{0,1}$ change linearly in \mathbb{R}^3 to the corresponding off-diagonal elements in $\mathbf{\Sigma}_{0,2}$.

Figure 2: Stationary model given by (7). The field looks the same wherever we are.

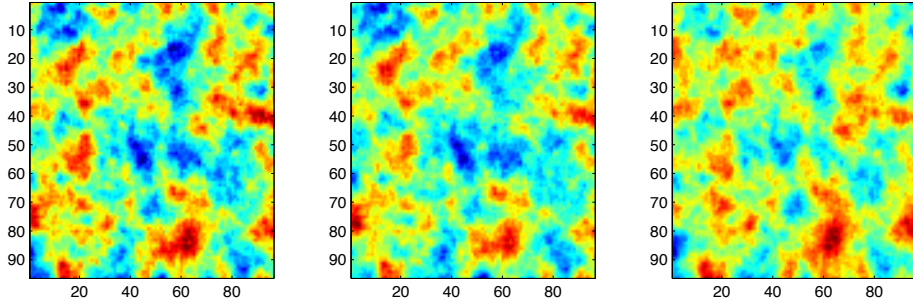


Figure 3: Nonstationary model with fixed \mathbf{Q}_0 . Here the bottom layer has anisotropy along the curve of the interface, and the correlation between the fields is fixed through space.

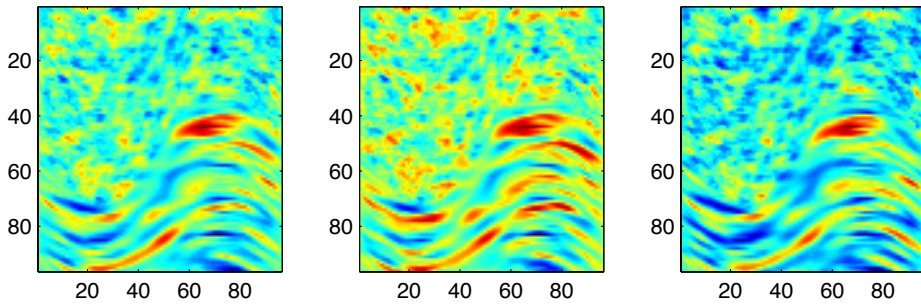
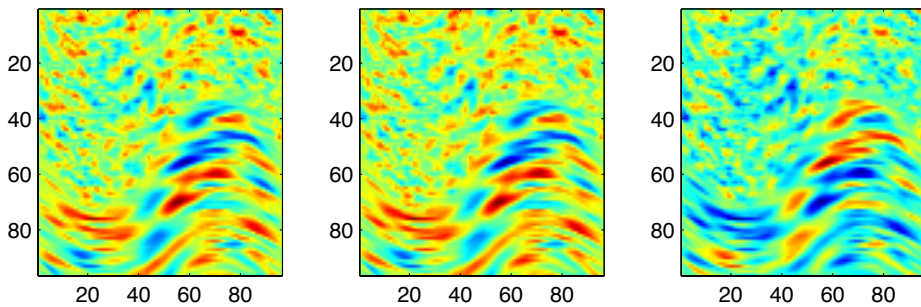


Figure 4: Full nonstationary model with varying $q_{ij}(\mathbf{s})$ according to (20). The bottom layer has anisotropy along the curved interface, and the correlation between the fields changes between interfaces. In particular, the rightmost field is positively correlated to the others above the interface and negatively correlated below the interface.



A very natural way of making such a transition between $\Sigma_{0,1}$ and $\Sigma_{0,2}$ is by considering geodesics on the manifold of symmetric positive definite matrices, denoted \mathbb{P}_d . The natural metric on this space has a reasonable statistical interpretation, closely related to information entropy and Kullback-Leibler divergence, and an accessible account for the theory is given in Bhatia (2007). Different treatments are given in (Ohara et al., 1996; Hiai and Petz, 2009). For completeness, we give a small account of the definition and properties we need related to this manifold. This exposition is based on Hiai and Petz (2009); Bhatia (2007).

The Boltzmann entropy of the Gaussian distribution (6), defining an information potential, is given by

$$B(p(\mathbf{x}|\mathbf{Q}, \boldsymbol{\mu}_x)) = B(\mathbf{Q}) = \frac{1}{2} \log \det \boldsymbol{\Sigma} + C, \quad (16)$$

where C is an arbitrary constant and $\boldsymbol{\Sigma} = \mathbf{Q}^{-1}$ is any positive definite matrix. The Riemannian metric based on this information potential is the Hessian

$$g_{\mathbf{Q}}(\mathbf{H}, \mathbf{M}) = \frac{\partial^2}{\partial s \partial t} \Bigg|_{s=0, t=0} B(\mathbf{Q} + s\mathbf{H} + t\mathbf{M}) = \text{tr } \mathbf{Q}\mathbf{H}\mathbf{Q}\mathbf{K}, \quad (17)$$

and where $\mathbf{H}, \mathbf{S} \in \mathbb{S}_d$, the tangent space of symmetric matrices, $\mathbb{S}_d = \{\mathbf{V} \in \mathbb{R}^{d \times d} | \mathbf{V} = \mathbf{V}^T\}$. This defines the line element

$$ds = \left(\text{tr} \left[(\mathbf{Q}^{-1/2} d\mathbf{Q} \mathbf{Q}^{-1/2})^2 \right] \right)^{1/2}. \quad (18)$$

Hence, if we have a curve in \mathbb{P}_d , i.e. $\gamma : [a, b] \rightarrow \mathbb{P}_d$, its length can be calculated as

$$L(\gamma) = \int_a^b \left(\text{tr} \left[(\gamma(t)^{1/2} \gamma'(t) \gamma(t)^{1/2})^2 \right] \right)^{1/2} dt \quad (19)$$

A nice property that follows from this is that lengths of curves are invariant under congruence transformations. That is, if $g(t) = \mathbf{X}^T \gamma(t) \mathbf{X}$, $L(\gamma) = L(g)$. The geodesic, the curve with minimal length, between two matrices, \mathbf{A} and \mathbf{B} can from this be deduced to be

$$g_{\mathbf{A}, \mathbf{B}}(t) = \mathbf{A} \#_t \mathbf{B} = \mathbf{A}^{1/2} \left(\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2} \right)^t \mathbf{A}^{1/2}, \quad t \in [0, 1]. \quad (20)$$

Obviously, $g_{\mathbf{A}, \mathbf{B}}(0) = \mathbf{A}$ and $g_{\mathbf{A}, \mathbf{B}}(1) = \mathbf{B}$. It is this curve we use when we go from $\mathbf{A} = \mathbf{Q}_{0,1} = \Sigma_{0,1}^{-1}$ to $\mathbf{B} = \mathbf{Q}_{0,2} = \Sigma_{0,2}$ in different discrete structures in our prior model, and this ensures that we are within the realm of positive definite matrices in a natural way. Noting that $(\mathbf{A} \#_t \mathbf{B})^{-1} = \mathbf{A}^{-1} \#_t \mathbf{B}^{-1}$, we see that it is unproblematic to work with precision matrices rather than covariance matrices. Integrating $g_{\mathbf{A}, \mathbf{B}}(t)$ yields the distance between the two matrices,

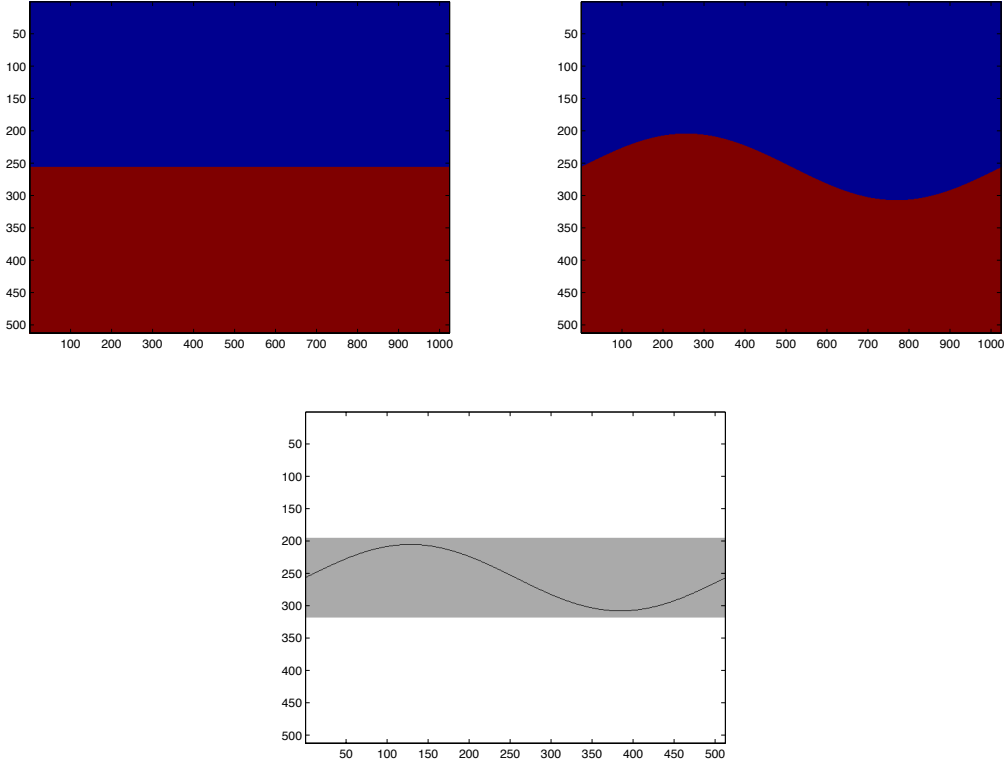
$$d_{\mathbb{P}_d}(\mathbf{A}, \mathbf{B}) = \int_0^1 g_{\mathbf{A}, \mathbf{B}}(t) = \left(\text{tr} \left[(\log \mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})^2 \right] \right)^{1/2}. \quad (21)$$

A potential drawback of using this strategy is that if $\mathbf{Q}_{0,1}, \mathbf{Q}_{0,2}$ are correlation matrices, and what you want is a continuum of correlation matrices, $g_{\mathbf{Q}_{0,1}, \mathbf{Q}_{0,2}}(t)$ are not correlation matrices for $t \in (0, 1)$. It is possible to correct for this by using geodesics on the submanifold of correlation matrices in \mathbb{P}_d . In practice, however, $g_{\mathbf{Q}_{0,1}, \mathbf{Q}_{0,2}}(t)$ are very close to being correlation matrices in most cases. We do not have any counterexamples.

3.2.1 Fuzzy interfaces

In some situations, we may actually have a hard interface in our multivariate field, but even in this situation, experts may place the interface incorrectly, which may lead to imprecise interpretation of the field. The geodesic blending strategy discussed in the previous section gives us a way to handle this situation in a specific way: the blending range may serve as

Figure 5: *Illustration of blend range. Gussed interface (left), true interface (right), blend range illustrated in grey (bottom)*



quantifying the uncertainty or fuzziness of this interface. This range may then be estimated based on realisations of the field, possibly requiring a strong prior for identifiability.

To illustrate this, suppose that an expert says that the interface is as in the upper left part of Figure 5, while the real interface is given on the right. The bottom illustration in Figure 5 shows what the geodesic range should be in this case (grey area) – it should cover the true interface properly, showing that there actually is a fair amount of uncertainty in the placement of the interface. In Section 4, we investigate whether this range may be estimated purely from data or if a strong prior on the range is needed. It is, of course, possible to combine this idea with procedures for actually estimating the interface, but, as always, this increases the complexity of the model that is to be estimated. Additionally, the blend range may easily confound with potential parameters needed to estimate the actual location of the interface.

3.3 Modified parametrisations

There are many ways to modify the parametrisation described above to reduce parametrisation demand or incorporate different flexibility. A possible way to reduce the parametrisation demand further is to do the modeling in the Cholesky domain. This is a simplification, but it is one we believe should increase interpretability and possibly estimation properties. To motivate this approach, consider the following: Suppose that the Cholesky factorisation of \mathbf{Q}_0 is given by $\mathbf{Q}_0 = \mathbf{L}_0 \mathbf{L}_0^T$, and that $\mathbf{Q}_{\text{space}} = \mathbf{B}_1^s (\mathbf{Q}_2^s)^*$, for some matrices $\mathbf{B}_1^s, \mathbf{B}_2^s$. Generating the matrices $\{\mathbf{B}_i^s\}_{1,2}$ can for instance be done by using $\alpha_s = \alpha/2$ in (8) and discretizing this operator, but there exist many other factorisations that may behave in better way for the problem at hand. By a Kronecker product identity, $\mathbf{Q}_{\text{space}} \otimes \mathbf{Q}_0 = (\mathbf{B}_1^s \otimes \mathbf{L}_0) ((\mathbf{B}_2^s)^* \otimes \mathbf{L}_0^T)$. The intuition stemming from this identity carries over to the more general case in a natural

way: Let $l_{ij}(\mathbf{s})$ be entry i, j of the Cholesky factor of the matrix $\{q_{ij}(\mathbf{s})\}_{ij}$ locally, and define locally operators that will correspond to some square root of its original form in (14). It is possible to define the operators in such a way that we get back (15), but this is of minor concern in practice as long as we get the interpretability we want. This is reminiscent to the triangular approach mentioned in Hu et al. (2012).

4 Parameter estimation and conditional expectation

In order to show that our proposed model is useful with confidence in the realm of seismic AVA inversion, we must show that estimation of hyper-parameters in the prior model is feasible and that the conditional expectation, $\mathbb{E}(\mathbf{m}|\mathbf{d}, \boldsymbol{\theta})$, is better than in the simpler model. A natural way to see if the hyper-parameters are identifiable is to simulate from the prior fields and do maximum likelihood estimates on these. If this works well, one may go one level higher and assume noisy observations of the form

$$\mathbf{y} = \mathbf{W}\mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}, \quad (22)$$

where \mathbf{W} denotes a convolution matrix defined by the wavelet, and \mathbf{A} denotes the reflectivity matrix, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. It is also here possible to do maximum likelihood estimates. For more information on estimating this type of model, consult your favourite treatise that discusses latent (Gaussian) models for, e.g. Rasmussen and Williams (2006); Rue and Held (2005); Cressie and Wikle (2011). In treating this estimation problem, we use the simpler anisotropic model where the correlation changes from positive to negative at an interface defined by a straight line. It is, of course, possible to estimate the geometry as well, but this is beyond the scope of this text.

When estimating the q_{ij} , supposing it changes between layers, we must impose constraints to enforce the interpretability we want – namely that of its local inverse being the correlation matrix of the multivariate field at that point. Now, the matrix consisting of ones on its diagonal with three free parameters off its diagonal uniquely specifies these constraints through its eigenvalues: they must all be greater than zero. Hence we have three constraints, depending only on the off-diagonal elements of the local correlation matrix. The same type of restriction would apply if we were to use general local 3×3 covariances instead. In that scenario, however, the three constraints would depend on six parameters instead of three. In this section, we will denote the different models as follows

1. Model 1 is the simple stationary $\mathbf{Q}_0 \otimes \mathbf{Q}_{\text{space}}$ as in (7)
2. Model 2 is stationary in space using the extended $q_{ij}(\mathbf{s})$ parametrisation as in Section 3.1, equation (15), using interpretability constraints
3. Model 3 is nonstationary in space and using the extended $q_{ij}(\mathbf{s})$ parametrisation as in Section 3.1, equation (15). Additionally, we use a blending of two correlation matrices at the interface, so that the correlation change is not discontinuous.

4.1 Identifiability

We show that the parameters in $\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1$ are identifiable through simulation. To do this, we simulate from many multivariate fields and estimate the parameters by maximum likelihood. If the estimated maximum likelihood density – which is estimated from several realisations – is unimodal, the parameters are identifiable. Suppose that $\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1$ are given by

$$\boldsymbol{\Sigma}_0 = \begin{pmatrix} 1 & 0.7 & 0.2 \\ & 1 & 0.4 \\ & & 1 \end{pmatrix}, \quad \boldsymbol{\Sigma}_1 = \begin{pmatrix} 1 & 0.7 & -0.9 \\ & 1 & -0.85 \\ & & 1 \end{pmatrix}, \quad (23)$$

using $\kappa^2 = 0.1$, and $\tau^2\mathbf{Q}$, with $\tau^2 = 50$. Using 200 realisations from the field, we get maximum likelihood density estimates for the parameters – these are illustrated in Figure 6. For obtaining the parameters, we used a quasi-Newton method with initial correlation parameters being zero. Judging from this figure, since all density estimates are unimodal, all parameters seem to be identifiable.

Figure 6: Maximum likelihood density estimates for correlation parameters, κ^2 and τ^2 using direct observations of 200 fields

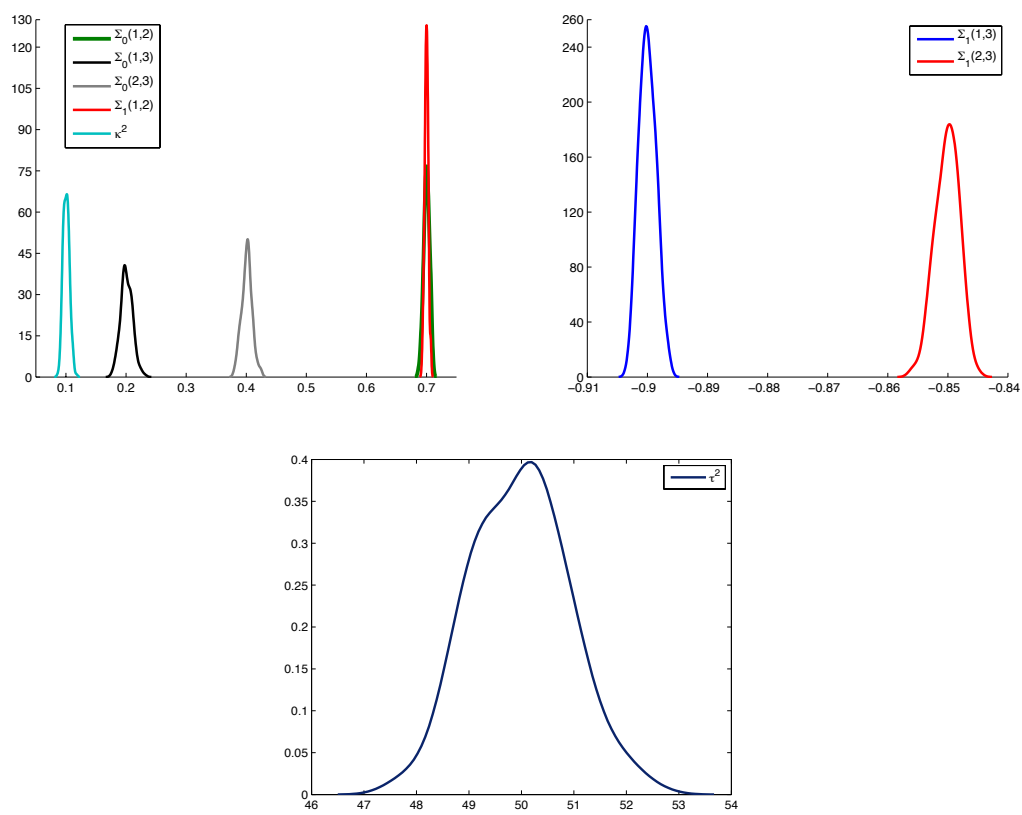
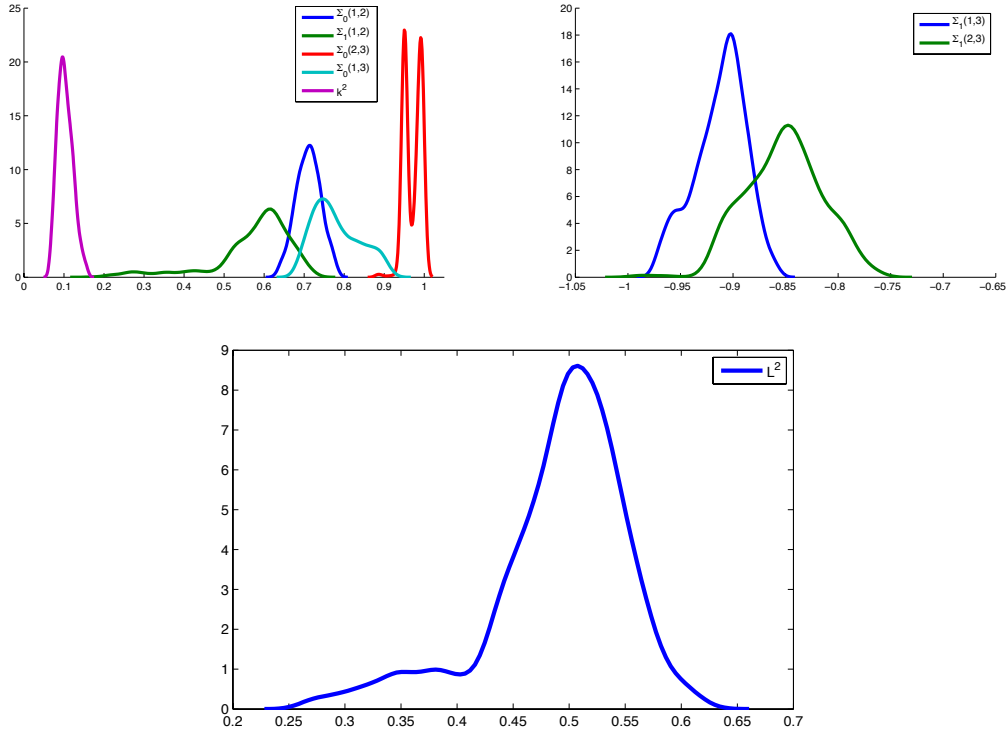


Figure 7: Maximum likelihood density estimates for correlation parameters, κ^2 and τ^2 using indirect observations of 600 fields



In the case where we have noisy observations, we use profile likelihood to estimate the noise level, σ^2 , and a quasi-Newton method to estimate $\lambda^2 = \tau^2\sigma^2$, κ^2 and the correlation parameters. In this case we used the correlation matrices

$$\Sigma_0 = \begin{pmatrix} 1 & 0.7 & 0.6 \\ & 1 & 0.95 \\ & & 1 \end{pmatrix}, \quad \Sigma_1 = \begin{pmatrix} 1 & 0.75 & -0.9 \\ & 1 & -0.85 \\ & & 1 \end{pmatrix}. \quad (24)$$

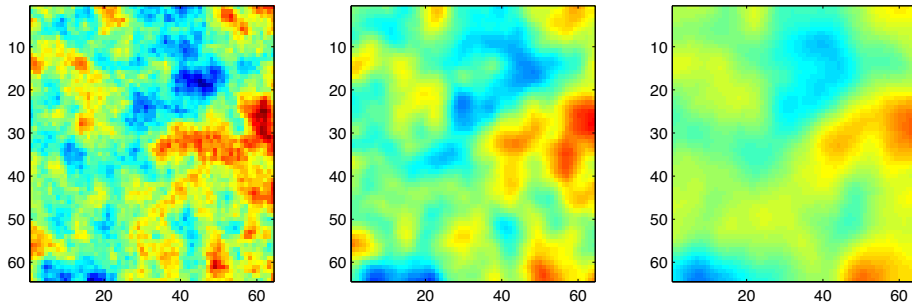
In Figure 7 the corresponding estimates for a hidden field is given. Of course, it is much more difficult in this situation, which is reflected through the broad distributional tails in the figure. Overall, however, the estimates seem to recover the true values quite well. One odd observation is the bimodality of $\Sigma_0(2,3)$. We believe it may come from observing rather small fields, from a 64×64 -grid, and that it may go away for larger ones. The values over one on the left part of the figure are artefacts coming from using a kernel smoother for estimating the density.

4.2 Conditional expectation

The real test on whether it is wise or not to use this advanced parametrisation of the model is essentially the reconstruction problem: based on noisy observations, are we able to reconstruct the original fields with higher fidelity?

In the following subsections, we give several reconstructions examples, and we use two observations schemes. The first one is based on identity observations with i.i.d. noise, and the second is based on the AVA model, giving the observation matrix \mathbf{WA} followed by i.i.d. noise.

Figure 8: *Kriging for identity observations with signal-to-noise ratio 1/50. True parameters (left), kriging using true Model 2 (center), using Model 1 (right)*



4.2.1 Identity observations

First, we present a reconstruction example where the observation matrix is the identity, followed by iid. noise, and with two signal-to-noise ratios. One with $\lambda^2 = 50$ and one with $\lambda^2 = 0.5$. For these two models, we use the following true correlation matrices

$$\Sigma_0 = \begin{pmatrix} 1 & 0.99 & 0.99 \\ & 1 & 0.99 \\ & & 1 \end{pmatrix}, \quad \Sigma_1 = \begin{pmatrix} 1 & -0.99 & -0.99 \\ & 1 & 0.99 \\ & & 1 \end{pmatrix}. \quad (25)$$

In Figure 8, we illustrate reconstruction of the first of the three fields with signal-to-noise ratio 1/50, using a flat interface and identity observations. I.e. the field true field is generated by Model 2, followed by i.i.d. noise. A priori we believe one of the worst situations for estimating Model 1, as correlations change very much from structure to structure and the noise level is very high. The likelihood function in the situation with high noise levels appears very flat, requiring high accuracy and many iterations in the optimisation scheme to give consistent estimates. For $\lambda^2 = 50$, $\|\mathbb{E}^{\text{M}2}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.526$ for Model 2 and $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.674$ for Model 1. The first field is chosen, as for the correlation matrices defined for this, the first field is the one with changing correlation between interfaces, relative to the others. The main effect we see in this comparison is that the level of the reconstructed field using the Model 1 does not completely reach up to the true levels – we believe this can be attributed to a flattening effect arising from the sum of the two correlations in the different layers being zero.

In the second comparison we generate fields and observations from Model 2, with $\lambda^2 = 0.5$. Here a different effect is more prominent – we see that the reconstructed field on the right in Figure 9, i.e. the reconstruction based on using Model 1, is smoother and does not exhibit as much of the jaggedy effect of the true surface compared to the field in the middle. A consistent property when estimating the Model 1 is that κ^2 seems to be underestimated, leading to a larger range and hence smoother reconstruction. One may think that this smoothing effect of the field on the right in Figure 9, but for comparison, we also include reconstructions of the second field, depicted in Figure 10. Here the mentioned smoothing effect is not as present as in Figure 9. Hence, we believe that this is an effect induced by the changing correlations. In Figure 9, $\|\mathbb{E}^{\text{M}2}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.179$ for the middle reconstruction, and $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.269$ for the rightmost one, while in Figure 10, $\|\mathbb{E}^{\text{M}2}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.168$ and $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.195$.

4.2.2 AVA observations

While the results using the identity observations are convincing in the extended models' favour, we also need to investigate the effects where the observation matrix is the seismic AVA model. In this case, the true fields are generated by Model 2, and the observations are

Figure 9: Kriging for identity observations with signal-to-noise ratio $1/0.5$, field 1. True parameters (left), kriging using Model 2 (center), using Model 1 (right)

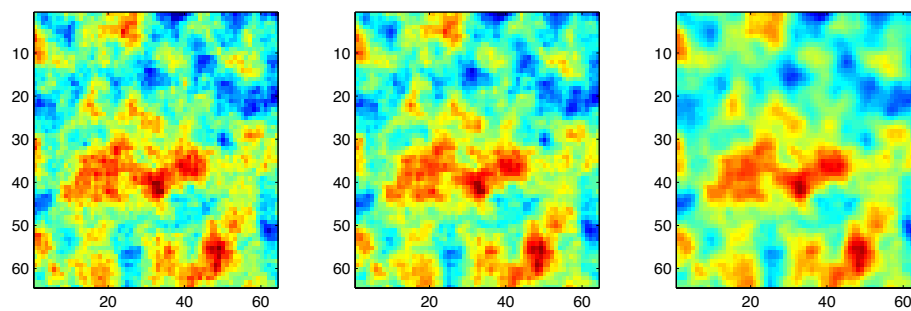


Figure 10: Kriging for identity observations with signal-to-noise ratio $1/0.5$, field 2. True parameters (left), using true model (center), using model defined by (12) (right)

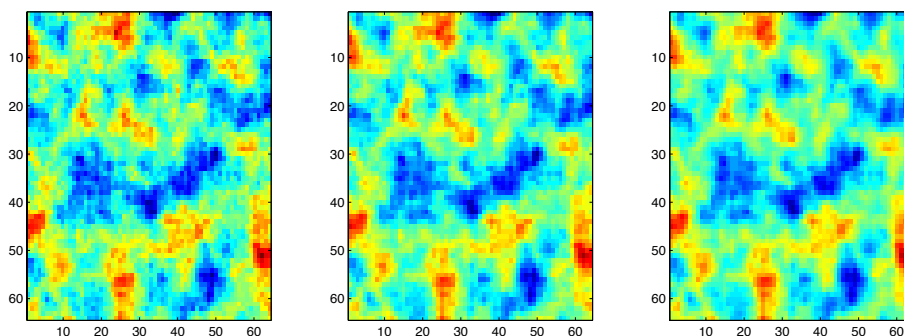
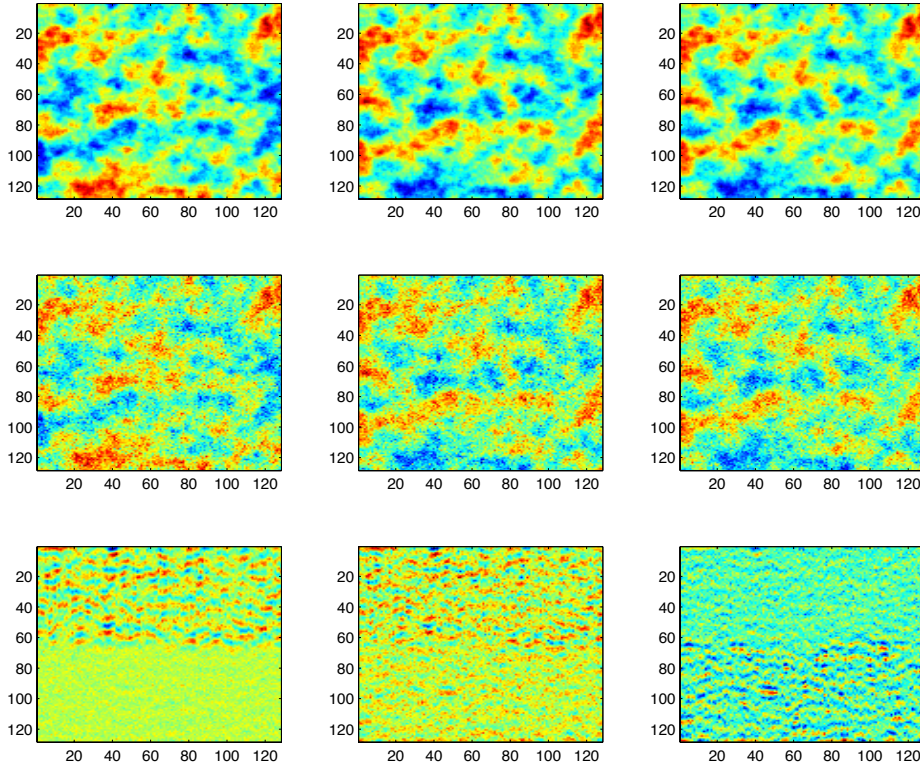


Figure 11: *Observations using identity observations (middle) and the seismic AVA model (bottom)*



linear combinations of the various fields at each space location followed by a convolution with a smooth wavelet and i.i.d noise. We use $\lambda^2 = 0.5$ and $\lambda^2 = 20$ in these examples.

In Figure 11, we can see the observations that are generated by this process. A key feature in the observations is that there occurs some cancellation, resulting from the fact that they are linear combinations of the underlying fields. This results in varying signal-to-noise ratios depending on the varying correlations.

Reconstructing the original multivariate field using the AVA observation scheme is more difficult than for using the identity observation matrix. The aforementioned cancellation effect is a major contributor to this. Additionally, there does not seem to be a straightforward way of interpreting the estimated correlation parameters coming from Model 1. In Figure 12, we illustrate the true parameters on the left, with reconstruction using the Model 2 the middle and the Model 1 on the right, using a signal-to-noise ratio of $1/0.5$. The effects we see are reminiscent of those using identity observations, but the smoothing effect is not present here. In this case $\|\mathbb{E}^{\text{M}2}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.461$, while $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.728$. Reconstruction using the same model, with a signal-to-noise ratio $1/20$ is depicted in Figure 13. No smoothing effect relative to Model 2 is observed here, but predictions are worse using Model 1, having $\|\mathbb{E}^{\text{M}2}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.762$, while $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.863$.

4.2.3 Identity observations and non-stationarity

Until this point, we have only studied the effects coming from changing correlations between interfaces. The model we have described is much richer than that, providing a flexible way of specifying anisotropy that moves along geometry of the subsurface. In this situation, we

Figure 12: Reconstructed field using the AVA model with signal-to-noise ratio 1/0.5. True parameters (left), kriging using Model 2 (centre), using Model 1 (right).

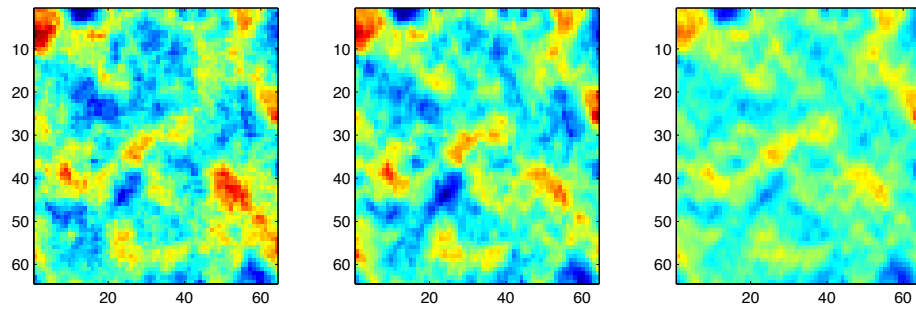


Figure 13: Reconstructed field using the AVA model with signal-to-noise ratio 1/20. True parameters (left), kriging using Model 2 (centre), using Model 1 (right).

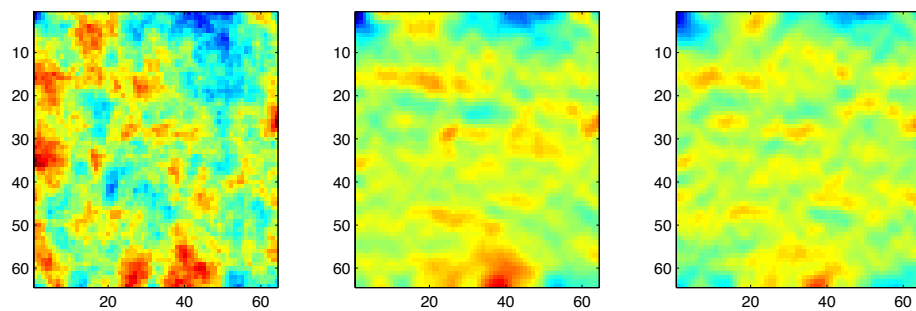
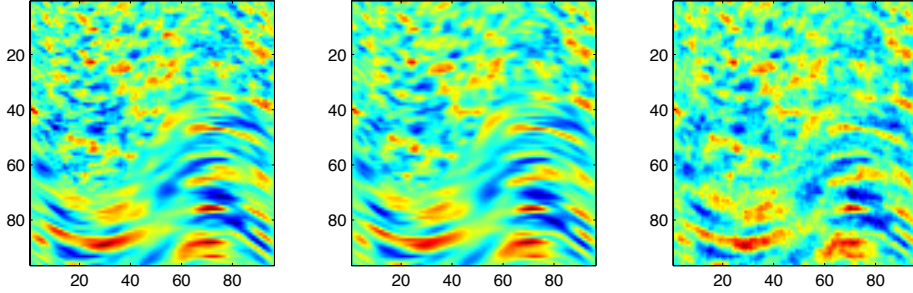


Figure 14: *Kriging for identity observations with signal-to-noise ration 1/0.2, field 1. True parameters (left), kriging using Model 3 (center), Model 1 (right)*



expect the results to be even more convincing, and we provide one example to cover this situation as well. In this case, we generate the true field by using Model 3, and we have identity observations followed by i.i.d. noise. The realisations of the true fields are then similar to the one in Figure 4, and we estimate both the simple and complex model for thereafter giving a reconstruction of the latent field. In Figure 14, we see the reconstructions using Model 3 (center) and Model 1 (right), and the most prominent effect we see is the smoothness differences in the bottom layer. Reconstruction using Model 1 is rugged and does not capture the anisotropy of the layer at all, contrasting the reconstruction using Model 3. On the top layer, on the other hand, the reconstructions are more comparable. The relative errors are $\|\mathbb{E}^{\text{M}3}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.280$ (center) and $\|\mathbb{E}^{\text{M}1}(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) - \mathbf{x}\|_2 / \|\mathbf{x}\|_2 = 0.382$ (right) for the reconstructions – i.e. predictions are about 37% better using the true model.

4.2.4 Reconstruction – final remarks

It is also important to note that if we simulate from the simple model, the parameters here are recovered well by using the parametrisation in Section 3.1. This means that the correlations estimated by the simple model are close to the ones estimated by the more complex model. This, of course, adds to the usefulness of the model in situations where we do not know in advance that the correlation changes between interfaces. The uncertainty, however, is greater, leading to more disparate estimates of the correlations than when using the simple parametrisation.

4.3 Estimating the blend range for fuzzy interfaces

In this section, we will treat all parameter except the blend range as fixed. The model we will treat is one where the true interface is given as a sine function, and what we guess is a flat interface. This is exactly the model which is depicted in Figure 5. In our example, we use Model 2 for constructing the true field, followed by identity observations and i.i.d. noise.

Before actually doing maximum likelihood estimation, we visualise heuristically why it may make sense. In Figure 15, we see a sample of the true model at the top, the true sample minus the guessed model in the middle, and the true sample minus the blend model with optimal range at the bottom. The norm of the bottom figure is less than that of the middle one.

Maximum likelihood estimates for the range is given in Figure 16, where the left figure is the range estimates when the guessed interface is a line and the true line is a full-period sine with a maximum amplitude of 23 and the right is a half-period sine with maximum amplitude 23. These estimates are good in the sense that the range covers the true model as in Figure 5.

A comparison of predictions using the guessed interface with no blend and the one with optimal blend is given in Figure 17. Here we see that the predictions using optimal blend

Figure 15: Sample from true model (top), sample from true minus guessed model (middle), sample from true model minus blend model with optimal range (bottom).

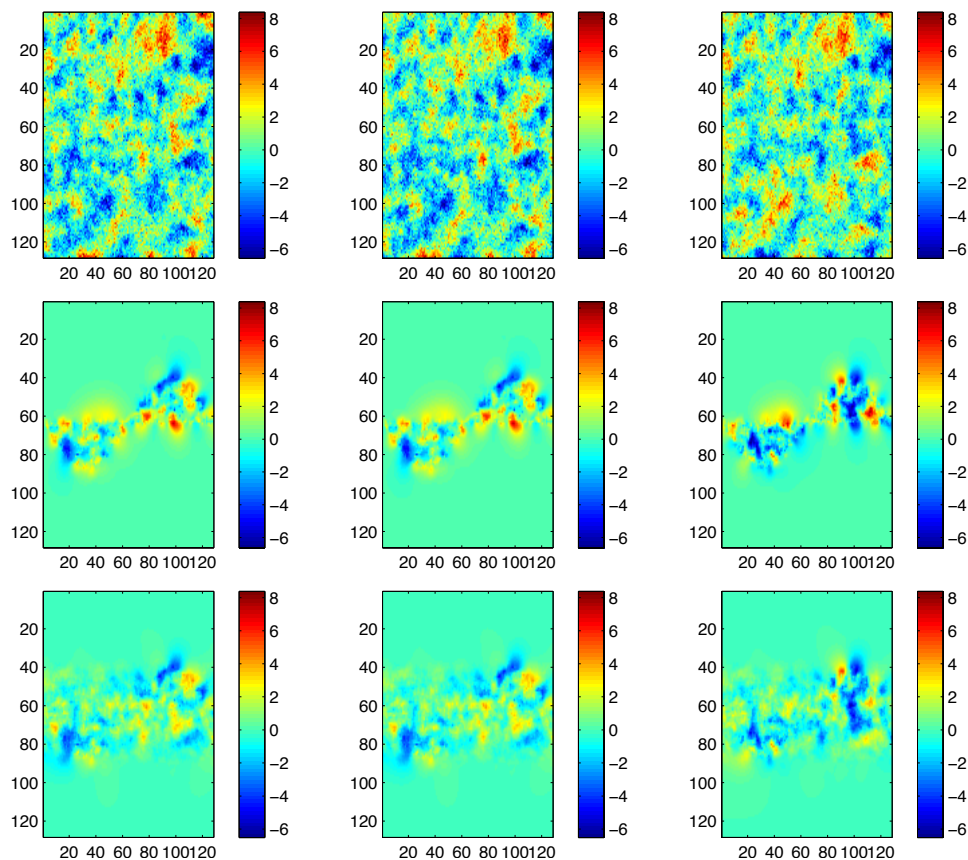


Figure 16: Estimated blend range using maximum likelihood for 200 samples. Sine-interface with full period (left), sine-interface with half period (right).

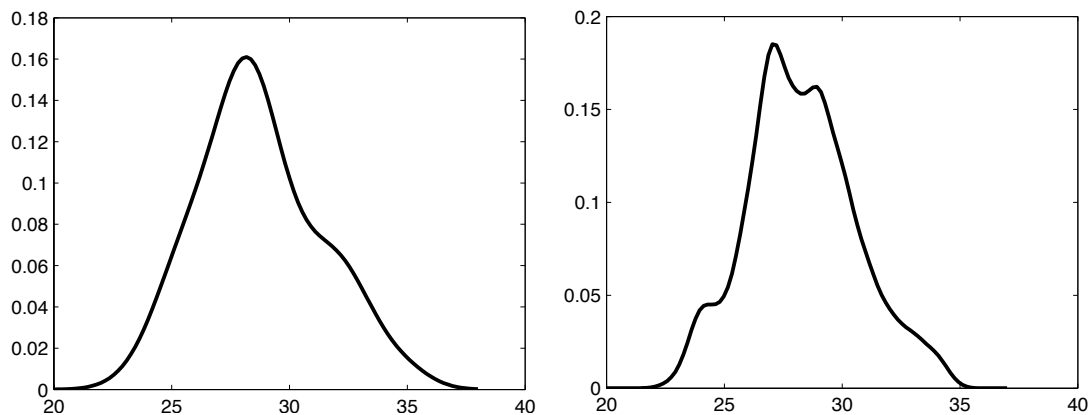
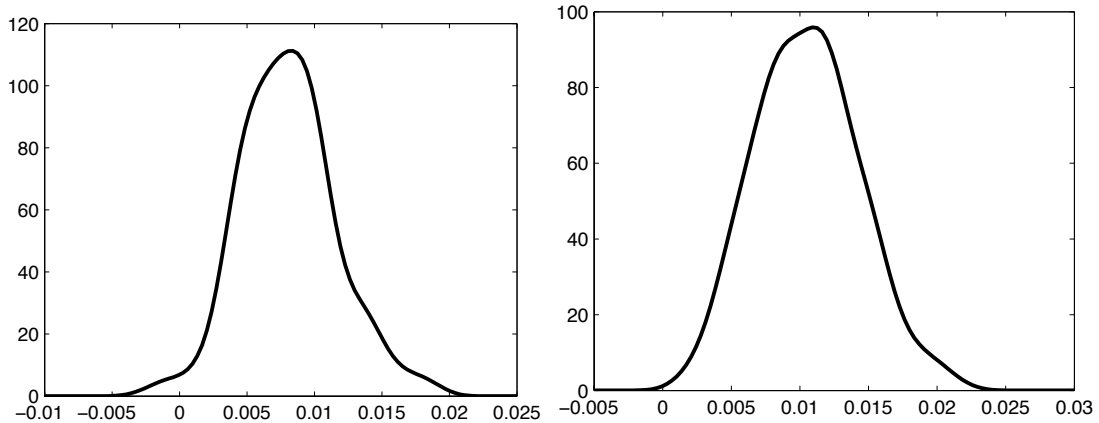


Figure 17: Performance gain using the blended interface over the guessed interface with sine-interface with full period (left) and sine-interface with half period (right)



range are only marginally better than using the interface with no blend for both interface structures. For the blend range, however, better prediction is not the goal. The goal here is to get an idea about how uncertain we are about the interface location, and better predictions comes as an additional boon, even if the improvement is marginal.

5 Conclusions and future work

In this text we have showed three things: First, how it is possible to incorporate information about the geometry of the problem flexibly. Secondly, how to facilitate changing covariances between elastic parameters depending on position. Lastly, we have introduced a novel way of specifying uncertainty related to the position of an interface using the concept of geodesic blending based on local correlation of the multivariate field. The first hinges on using SPDEs in order to specify local properties of the fields, and the second on how systems of SPDEs interrelate depending on position. The geodesic blending approach is based on the smooth manifold structure of the set of positive definite matrices. The ideas presented here are not limited to the relatively simple models described here – rather, they may be used in any spatial inversion problem with a natural geometry where soft constraints based on expert opinion may be used.

Address for corresponding author:

Erlend Aune
 Nedre Møllenberggate 70B
 7043 Trondheim
 E-mail: erlend.aune.1983@gmail.com

Appendix: Finite difference discretization – the gory details

This appendix is devoted to the finite difference scheme we used for discretizing the elliptic operator in (11). We employ a changed notation in this appendix for convenience, replacing \mathbf{H} with \mathbf{A} , and we hope that it is transparent for readers. For a 2-dimensional field with $\alpha = 1$, we have

$$\begin{aligned}
& \nabla \cdot \begin{pmatrix} a_{11}(x, y) & a_{12}(x, y) \\ a_{21}(x, y) & a_{22}(x, y) \end{pmatrix} \begin{pmatrix} u_x(x, y) \\ u_y(x, y) \end{pmatrix} + \kappa(x, y)u(x, y) \\
&= \nabla \cdot \begin{pmatrix} a_{11}(x, y)u_x(x, y) + a_{1,2}(x, y)u_y(x, y) \\ a_{21}(x, y)u_x(x, y) + a_{22}(x, y)u_y(x, y) \end{pmatrix} + \kappa(x, y)u(x, y) \\
&= \partial_x(a_{11}u_x + a_{12}u_y) + \partial_y(a_{21}u_x + a_{22}u_y) + \kappa u \\
&= a_{11}^x u_x + a_{11}u_{xx} + a_{12}^x u_y + a_{12}u_{yx} + a_{21}^y u_x + a_{21}u_{xy} + a_{22}^y u_y + a_{22}u_{yy} \\
&= \text{diag}(A)\nabla \cdot \nabla u + (a_{12} + a_{21})u_{xy} + a_{11}^x u_x + a_{12}^x u_y + a_{22}^y u_y + a_{21}^y u_x
\end{aligned} \tag{26}$$

where a_{ij}^v , $v = x, y$ denotes differentiation wrt. x or y of the i, j element of \mathbf{A} , depending implicitly on the position. To discretize (26), we employ a finite difference scheme. We define the following finite difference operators

$$\begin{aligned}
\delta_x u &= \frac{1}{h}(u_{i+1}^j - u_i^j) \\
\delta_{\hat{x}} &= \frac{1}{h}(u_i^j - u_{i-1}^j),
\end{aligned}$$

where i, j are positions on the grid, with i denoting the x -direction and j denoting the y -direction. Now, we define the following operators

$$\begin{aligned}
\Lambda_{xx} u &= \delta_x (\alpha_{11} \delta_{\hat{x}} u) = \delta_x \left(\frac{1}{h} \alpha_{11} (u_i^j - u_{i-1}^j) \right) \\
&= \frac{1}{h^2} \left(\alpha_{11}^{i+1, j} (u_{i+1}^j - u_i^j) - \alpha_{11}^{i, j} (u_i^j - u_{i-1}^j) \right),
\end{aligned} \tag{27}$$

where

$$\begin{aligned}
\alpha_{11}^{i, j} &= \frac{1}{2} (a_{11}^{i, j} + a_{11}^{i-1, j}) \\
\alpha_{22}^{i, j} &= \frac{1}{2} (a_{22}^{i, j} + a_{22}^{i, j-1}).
\end{aligned}$$

A equivalent expression holds for $\Lambda_{yy} u$. We define $\alpha_{kk}^{1,1} = a_{11}^{1,1}$, $k = 1, 2$. For the mixed operators we have

$$\Lambda_{xy}^+ u = \frac{1}{2} (\delta_x (a_{12} \delta_y u) + \delta_{\hat{x}} (a_{12} \delta_{\hat{y}} u)) \tag{28}$$

and we have

$$\begin{aligned}
\delta_x (a_{12} \delta_y u) &= \frac{1}{h} \delta_x (a_{12} u_i^{j+1} - u_i^j) \\
&= \frac{1}{h^2} (a_{12}^{i+1, j} (u_{i+1}^{j+1} - u_{i+1}^j) - a_{12}^{i, j} (u_i^{j+1} - u_i^j)) \\
\delta_{\hat{x}} (a_{12} \delta_{\hat{y}} u) &= \frac{1}{h} \delta_{\hat{x}} (a_{12} (u_i^j - u_i^{j-1})) \\
&= \frac{1}{h^2} (a_{12}^{i, j} (u_i^j - u_i^{j-1}) - a_{12}^{i-1, j} (u_{i-1}^j - u_{i-1}^{j-1})).
\end{aligned}$$

Hence

$$\begin{aligned}
\Lambda_{xy}^+ u &= \frac{1}{2h^2} \left((a_{12}^{i+1, j} (u_{i+1}^{j+1} - u_{i+1}^j) - a_{12}^{i, j} (u_i^{j+1} - u_i^j)) \right. \\
&\quad \left. + (a_{12}^{i, j} (u_i^j - u_i^{j-1}) - a_{12}^{i-1, j} (u_{i-1}^j - u_{i-1}^{j-1})) \right)
\end{aligned} \tag{29}$$

For Λ_{yx}^+ we reverse the order of the difference operators:

$$\begin{aligned}\Lambda_{yx}^+ u &= \frac{1}{2}(\delta_y(a_{12}\delta_x u) + \delta_{\hat{y}}(a_{12}\delta_{\hat{x}})) \\ &= \frac{1}{2h^2} \left(\left(a_{12}^{i,j+1}(u_{i+1}^{j+1} - u_i^{j+1}) - a_{12}^{i,j}(u_{i+1}^j - u_i^j) \right) \right. \\ &\quad \left. + \left(a_{12}^{i,j}(u_i^j - u_{i-1}^j) - a_{12}^{i,j-1}(u_i^{j-1} - u_{i-1}^{j-1}) \right) \right)\end{aligned}$$

And the complete discretisation is

$$(\Lambda_{xx} + \Lambda_{xy}^+ + \Lambda_{yx}^+ + \Lambda_{yy})u = f(u, W) \quad (30)$$

In Samarskii et al. (2002), it is proved that this scheme is convergent. If we assume that \mathbf{A} does not vary in space, we can simplify the scheme;

$$\begin{aligned}\widehat{\Lambda_{xx}} u &= \frac{1}{h^2}(a_{11}(u_{i+1}^j - u_i^j) - a_{11}(u_i^j - u_{i-1}^j)) \\ &= \frac{1}{h^2}a_{11}(u_{i+1}^j - 2u_i^j + u_{i-1}^j) \\ \widehat{\Lambda_{yy}} u &= \frac{1}{h^2}a_{22}(u_i^{j+1} - 2u_i^j + u_i^{j-1}) \\ \widehat{\Lambda_{xy}^+} u &= \frac{1}{2h^2} \left(\left(a_{12}(u_{i+1}^{j+1} - u_{i+1}^j) - a_{12}(u_i^{j+1} - u_i^j) \right) \right. \\ &\quad \left. + \left(a_{12}(u_i^j - u_{i-1}^{j-1}) - a_{12}(u_{i-1}^j - u_{i-1}^{j-1}) \right) \right) \\ &= \frac{a_{12}}{2h^2} \left(2u_i^j + u_{i+1}^{j+1} + u_{i-1}^{j-1} - u_{i+1}^j - u_i^{j+1} - u_i^{j-1} - u_{i-1}^j \right) \\ \widehat{\Lambda_{yx}^+} u &= \frac{1}{2h^2} \left(\left(a_{12}(u_{i+1}^{j+1} - u_i^{j+1}) - a_{12}(u_{i+1}^j - u_i^j) \right) \right. \\ &\quad \left. + \left(a_{12}(u_i^j - u_{i-1}^j) - a_{12}(u_i^{j-1} - u_{i-1}^{j-1}) \right) \right) \\ &= \frac{a_{12}}{2h^2} \left(2u_i^j + u_{i+1}^{j+1} + u_{i-1}^{j-1} - u_{i+1}^j - u_i^{j+1} - u_i^{j-1} - u_{i-1}^j \right) \\ \left(\widehat{\Lambda_{xy}^+} + \widehat{\Lambda_{yx}^+} \right) u &= \frac{a_{12}}{h^2} \left(2u_i^j + u_{i+1}^{j+1} + u_{i-1}^{j-1} - u_{i+1}^j - u_i^{j+1} - u_i^{j-1} - u_{i-1}^j \right)\end{aligned}$$

This corresponds to the following stencil

$$S = -\frac{1}{h^2} \begin{pmatrix} a_{12} & -a_{22} - a_{12} & 0 \\ -a_{11} - a_{12} & 2(a_{11} + a_{22} + a_{12}) & -a_{11} - a_{12} \\ 0 & -a_{22} - a_{12} & a_{12} \end{pmatrix} \quad (31)$$

References

- Aki, K. and Richards, P. G. (1980). *Quantitative Seismology*. WH Freeman and Co.
- Aune, E., Eidsvik, J., and Pokern, Y. (2012a). Iterative Numerical Methods for Sampling from High Dimensional Gaussian Distributions. *Statistics and Computing*, Accepted.
- Aune, E., Simpson, D. P., and Eidsvik, J. (2012b). Parameter estimation for high dimensional Gaussian distributions. *Statistics and Computing. Revised*.
- Bhatia, R. (2007). *Positive definite matrices*. Princeton Univ Press.
- Buland, A., Kolbjørnsen, O., and Omre, H. (2003). Rapid spatially coupled AVO inversion in the Fourier domain. *Geophysics*, 68:824–836.
- Buland, A. and Omre, H. (2003). Bayesian linearized AVO inversion. *Geophysics*, 68:185–198.
- Cressie, N. (1993). *Statistics for Spatial Data*. Wiley.
- Cressie, N. and Wikle, C. (2011). *Statistics for spatio-temporal data*, volume 465. Wiley.
- Feichtinger, H., Helffer, B., Lamoureux, M., Lerner, N., Morel, J., Rodino, L., Takens, F., Teissier, B., Toft, J., and Wong, M. (2008). *Pseudo-Differential Operators: Quantization and Signals*. Springer-Verlag Berlin Heidelberg.
- Gneiting, T., Kleiber, W., and Schlather, M. (2010). Matérn cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association*, 105(491):1167–1177.
- Gray, R. (2006). *Toeplitz and circulant matrices: A review*. E-book.
- Hiai, F. and Petz, D. (2009). Riemannian metrics on positive definite matrices related to means. *Linear Algebra and Its Applications*, 430(11-12):3105–3130.
- Hu, X., Simpson, D. P., Lindgren, F., and Rue, H. (2012). Multivariate gaussian random fields using stochastic partial differential equations. *N/A*.
- Lindgren, F., Lindstrøm, J., and Rue, H. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The SPDE approach. *Journal of the Royal Statistical Society, Series B*, 73:423–498.
- Matérn, B. (1960). Spatial variation. *Meddelanden fran Statens Skogsforskningsinstitut (Stockholm)*, 49.
- Ohara, A., Suda, N., and Amari, S. (1996). Dualistic differential geometry of positive definite matrices and its applications to related problems. *Linear Algebra and Its Applications*, 247:31–53.
- Rabben, T. E., Ursin, B., and Tjelmeland, H. (2008). Non-linear Bayesian joint inversion of seismic reflection coefficients. *Geophysical journal international*, 173:265–280.
- Rasmussen, C. and Williams, C. (2006). *Gaussian processes for machine learning*. MIT Press, MA.
- Rozanov, J. (1977). Markov random fields and stochastic partial differential equations. *Mathematics of the USSR-Sbornik*, 32:515.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields*. Chapman & Hall.
- Ruzhansky, M. and Turunen, V. (2009). *Pseudo-differential Operators and Symmetries: Background Analysis and Advanced Topics*. Birkhauser.

- Samarskii, A. A., Matus, P. P., Mazhukin, V. I., and Mozolevski, I. E. (2002). Monotone Difference Schemes for Equations with Mixed Derivatives. *Computers and Mathematics*, 44:501 – 510.
- Simpson, D. (2008). *Krylov subspace methods for approximating functions of symmetric positive definite matrices with applications to applied statistics and anomalous diffusion*. PhD thesis, School of Mathematical Sciences, Queensland Univ of Tech.
- Stovas, A. and Ursin, B. (2003). Reflection and transmission responses of layered transversely isotropic viscoelastic media. *Geophysical Prospecting*, 51:447–477.
- Whittle, P. (1963). Stochastic processes in several dimensions. *Bull. Inst. Internat. Statist*, 40:974–994.