

Stochastic partial differential equations and numerical methods for large scale spatial statistics

Finn Lindgren



Big Data in Environmental Science
University of British Columbia, PIMS
Vancouver, 11 May 2015

“Big” data

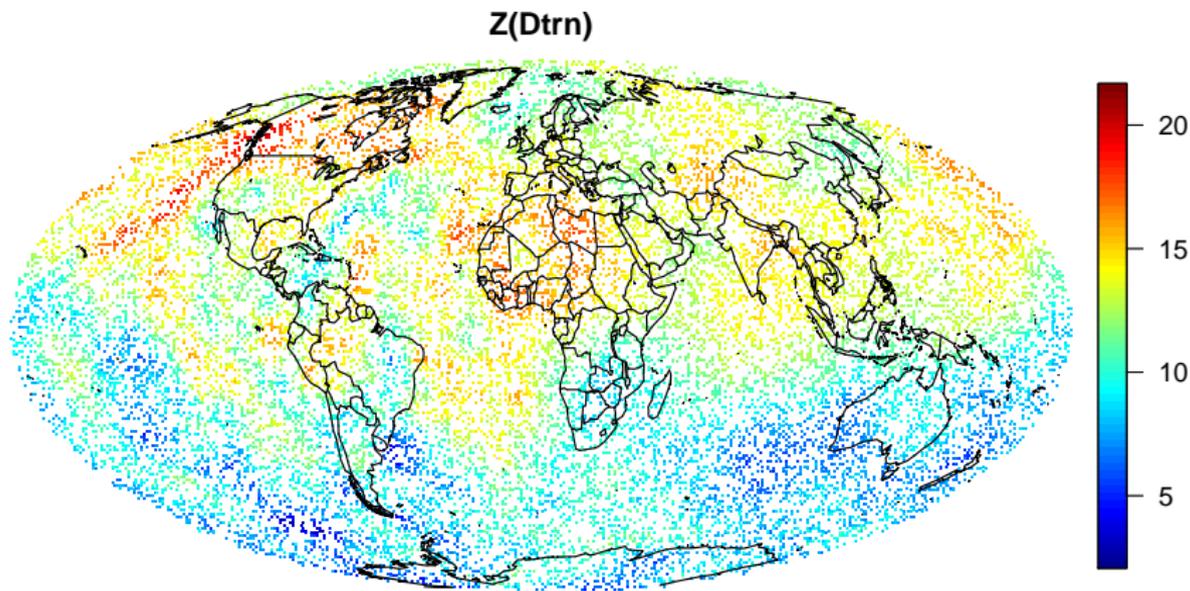
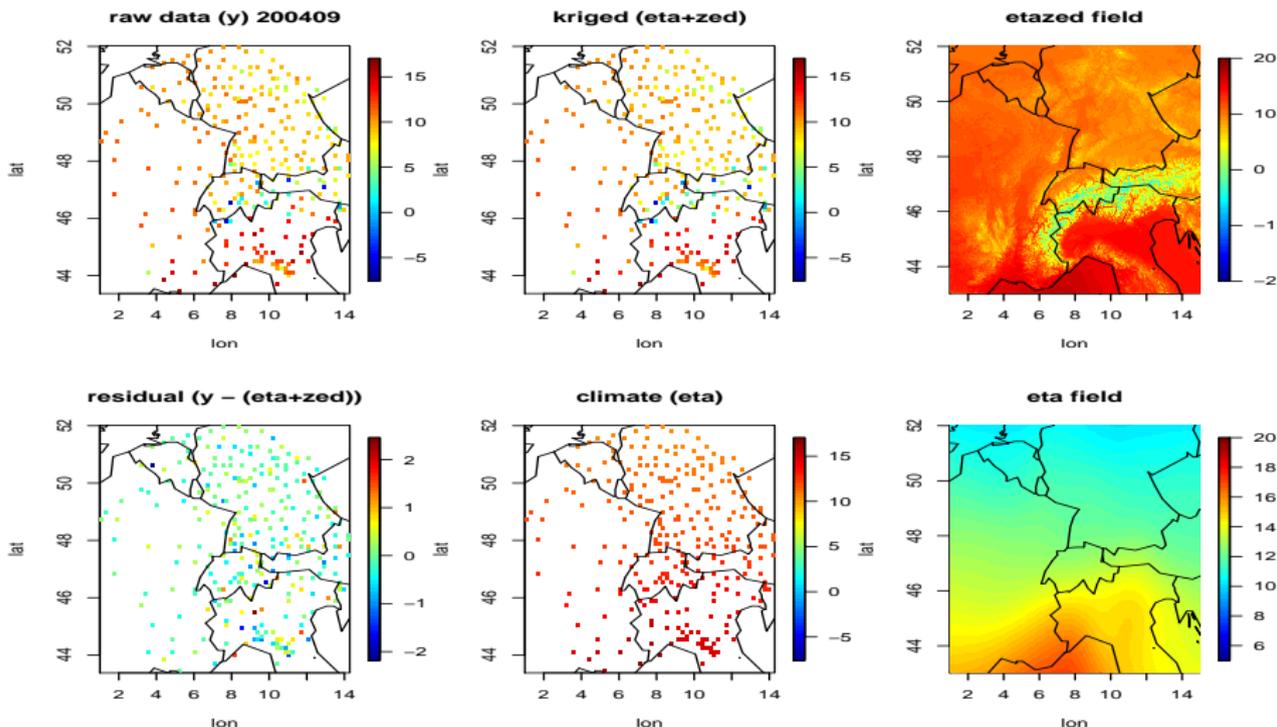


Illustration: Synthetic data mimicking satellite based CO_2 measurements. Irregular data locations, uneven coverage, features on different scales.

Sparse spatial coverage of temperature measurements



Regional observations: $\approx 20,000,000$ from daily timeseries over 160 years

Note: This is a small *subset* of the full data! (More about this tomorrow!)

Spatio-temporal modelling framework

Spatial statistics framework

- ▶ Spatial domain Ω , or space-time domain $\Omega \times \mathbb{T}$, $\mathbb{T} \subset \mathbb{R}$.
- ▶ Random field $u(\mathbf{s})$, $\mathbf{s} \in \Omega$, or $u(\mathbf{s}, t)$, $(\mathbf{s}, t) \in \Omega \times \mathbb{T}$.
- ▶ Observations y_i . In the simplest setting, $y_i = u(\mathbf{s}_i) + \epsilon_i$, but more generally $y_i \sim \text{GLMM}$, with $u(\cdot)$ as a structured random effect.
- ▶ Needed: models capturing stochastic dependence on multiple scales
- ▶ Partial solution: Basis function expansions, with large scale functions and covariates to capture static and slow structures, and small scale functions for more local variability

Two basic model and method components

- ▶ Stochastic models for $u(\cdot)$.
- ▶ Computationally efficient (i.e. avoid MCMC whenever possible) inference methods for the posterior distribution of $u(\cdot)$ given data \mathbf{y} .

Covariance functions and stochastic PDEs

The Matérn covariance family on \mathbb{R}^d

$$\text{Cov}(u(\mathbf{0}), u(\mathbf{s})) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} (\kappa \|\mathbf{s}\|)^\nu K_\nu(\kappa \|\mathbf{s}\|)$$

Scale $\kappa > 0$, smoothness $\nu > 0$, variance $\sigma^2 > 0$



Whittle (1954, 1963): Matérn as SPDE solution

Matérn fields are the stationary solutions to the SPDE

$$(\kappa^2 - \nabla \cdot \nabla)^{\alpha/2} u(\mathbf{s}) = \mathcal{W}(\mathbf{s}), \quad \alpha = \nu + d/2$$

$\mathcal{W}(\cdot)$ white noise, $\nabla \cdot \nabla = \sum_{i=1}^d \frac{\partial^2}{\partial s_i^2}$, $\sigma^2 = \frac{\Gamma(\nu)}{\Gamma(\alpha) \kappa^{2\nu} (4\pi)^{d/2}}$

(Continuous domain white noise, $\mathbb{E}(\mathcal{W}(A)) = 0$, and

$\text{Cov}(\mathcal{W}(A), \mathcal{W}(B)) = |A \cap B|$, for all measurable $A, B \subseteq \mathbb{R}^d$. Not to be confused with pointwise independent noise.)



Basis function representations for Gaussian Matérn fields

Basis definitions

	Finite basis set ($k = 1, \dots, n$)
Karhunen-Loève	$(\kappa^2 - \nabla \cdot \nabla)^{-\alpha} e_{\kappa,k}(\mathbf{s}) = \lambda_{\kappa,k} e_{\kappa,k}(\mathbf{s})$
Fourier	$-\nabla \cdot \nabla e_k(\mathbf{s}) = \lambda_k e_k(\mathbf{s})$
Convolution	$(\kappa^2 - \nabla \cdot \nabla)^{\alpha/2} g_\kappa(\mathbf{s}) = \delta(\mathbf{s})$
General/GMRF	$\psi_k(\mathbf{s})$

Field representations

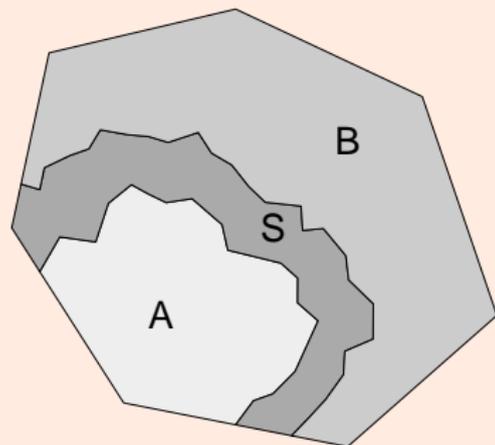
	Field $u(\mathbf{s})$	Weights
Karhunen-Loève	$\propto \sum_k e_{\kappa,k}(\mathbf{s}) z_k$	$z_k \sim \mathcal{N}(0, \lambda_{\kappa,k})$
Fourier	$\propto \sum_k e_k(\mathbf{s}) z_k$	$z_k \sim \mathcal{N}(0, (\kappa^2 + \lambda_k)^{-\alpha})$
Convolution	$\propto \sum_k g_\kappa(\mathbf{s} - \mathbf{s}_k) z_k$	$z_k \sim \mathcal{N}(0, \text{cell}_k)$
General/GMRF	$\propto \sum_k \psi_k(\mathbf{s}) u_k$	$\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_\kappa^{-1})$

Note: Harmonic basis functions (as in the Fourier approach) give a diagonal \mathbf{Q}_κ , but lead to dense posterior precision matrices.

Continuous and discrete Markov properties

Markov properties

S is a separating set for A and B : $u(A) \perp u(B) \mid u(S)$



Solutions to

$$(\kappa^2 - \nabla \cdot \nabla)^{\alpha/2} u(s) = \mathcal{W}(s)$$

are Markov when α is an integer.

(Generally, when the reciprocal of the spectral density is a polynomial, Rozanov, 1977)

Discrete representations ($Q = \Sigma^{-1}$):

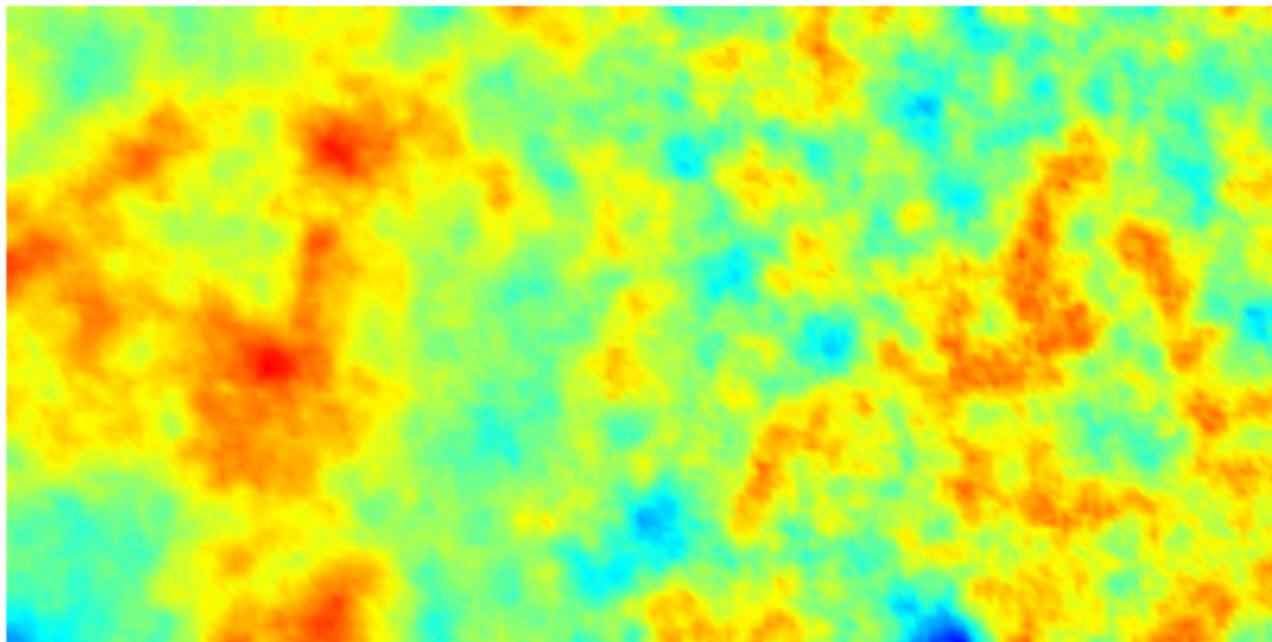
$$Q_{AB} = 0$$

$$Q_{A|S,B} = Q_{AA}$$

$$\mu_{A|S,B} = \mu_A - Q_{AA}^{-1} Q_{AS} (u_S - \mu_S)$$

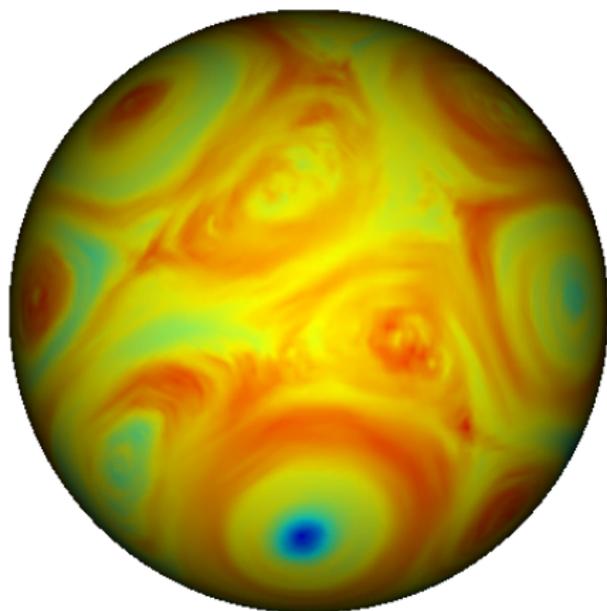
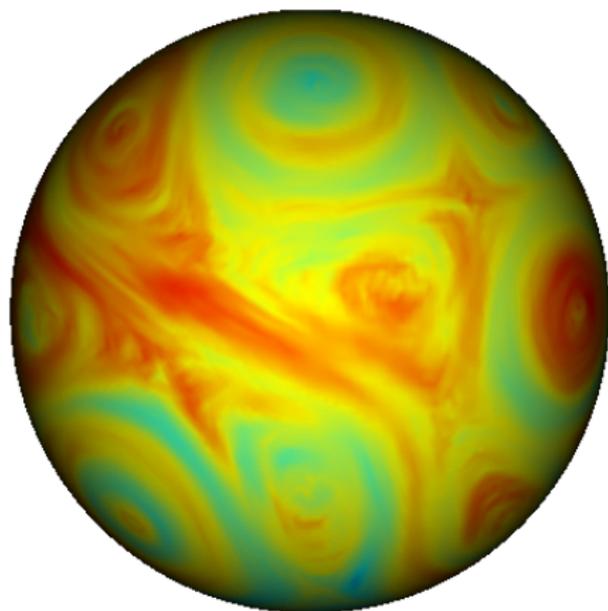
Fractional SPDEs can often be closely approximated by sums of a small number of Markov SPDEs, or even by a single (non-Matérn) Markov SPDE.

Beyond Matérn: Non-stationary field



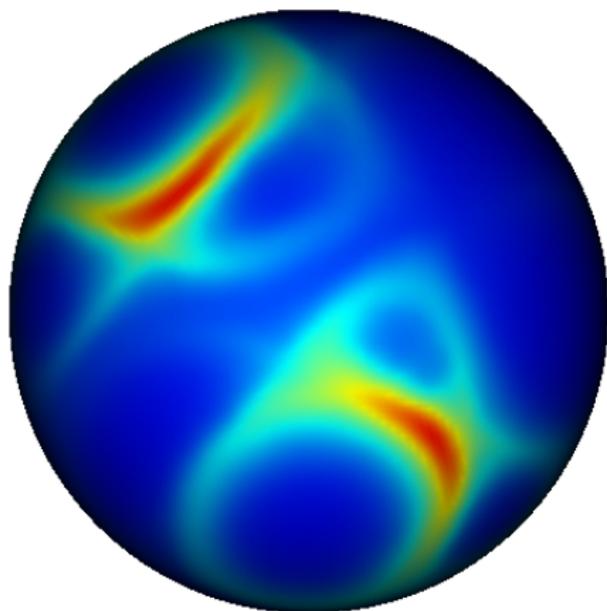
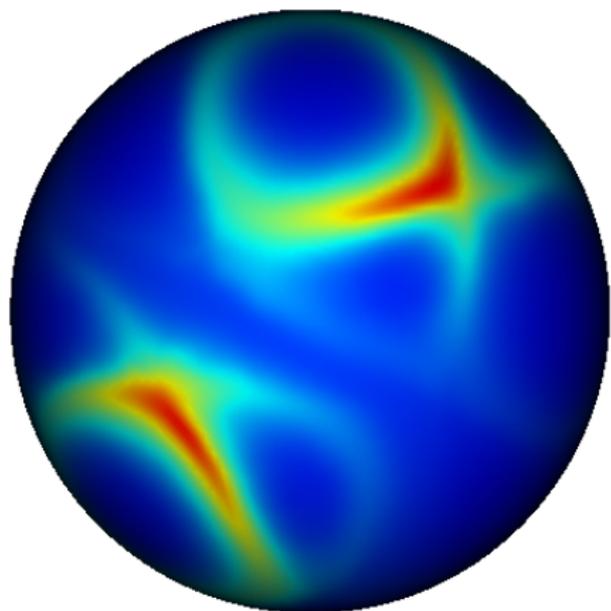
$$(\kappa(\mathbf{s}))^2 - \nabla \cdot \nabla)u(\mathbf{s}) = \kappa(\mathbf{s})\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \Omega$$

Beyond Matérn: Markov does *not* mean local dependence



$$(\kappa(\mathbf{s})^2 - \nabla \cdot \mathbf{H}(\mathbf{s})\nabla)u(\mathbf{s}) = \kappa(\mathbf{s})\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \Omega$$

Covariances for four reference points



Continuous domain Markov approximations

Continuous Markovian spatial models (Lindgren et al, 2011)

Local basis: $u(\mathbf{s}) = \sum_k \psi_k(\mathbf{s}) u_k$, (compact, piecewise linear)

Basis weights: $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1})$, sparse \mathbf{Q} based on an SPDE

Special case: $(\kappa^2 - \nabla \cdot \nabla)u(\mathbf{s}) = \mathcal{W}(\mathbf{s})$, $\mathbf{s} \in \Omega$

Precision: $\mathbf{Q} = \kappa^4 \mathbf{C} + 2\kappa^2 \mathbf{G} + \mathbf{G}_2$ ($\kappa^4 + 2\kappa^2|\boldsymbol{\omega}|^2 + |\boldsymbol{\omega}|^4$)

Conditional distribution in a jointly Gaussian model

$$\mathbf{u} \sim \mathcal{N}(\boldsymbol{\mu}_u, \mathbf{Q}_u^{-1}), \quad \mathbf{y}|\mathbf{u} \sim \mathcal{N}(\mathbf{A}\mathbf{u}, \mathbf{Q}_{y|\mathbf{u}}^{-1}) \quad (A_{ij} = \psi_j(\mathbf{s}_i))$$

$$\mathbf{u}|\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{u|\mathbf{y}}, \mathbf{Q}_{u|\mathbf{y}}^{-1})$$

$$\mathbf{Q}_{u|\mathbf{y}} = \mathbf{Q}_u + \mathbf{A}^T \mathbf{Q}_{y|\mathbf{u}} \mathbf{A} \quad (\sim \text{Sparse iff } \psi_k \text{ have compact support})$$

$$\boldsymbol{\mu}_{u|\mathbf{y}} = \boldsymbol{\mu}_u + \mathbf{Q}_{u|\mathbf{y}}^{-1} \mathbf{A}^T \mathbf{Q}_{y|\mathbf{u}} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_u)$$

We've translated the spatial inference problem into sparse numerical linear algebra similar to finite element PDE solvers

The computational GMRF work-horse

Cholesky decomposition (Cholesky, 1924)

$$Q = LL^T, \quad L \text{ lower triangular } (\sim \mathcal{O}(n^{(d+1)/2}) \text{ for } d = 1, 2, 3)$$

$$Q^{-1}\mathbf{x} = L^{-T}L^{-1}\mathbf{x}, \quad \text{via forward/backward substitution}$$

$$\log \det Q = 2 \log \det L = 2 \sum_i \log L_{ii}$$

André-Louis Cholesky (1875–1918)

"He invented, for the solution of the condition equations in the method of least squares, a very ingenious computational procedure which immediately proved extremely useful, and which most assuredly would have great benefits for all geodesists, if it were published some day." (Euology by Commandant Benoit, 1922)



Laplace approximations for non-Gaussian observations

Quadratic posterior log-likelihood approximation

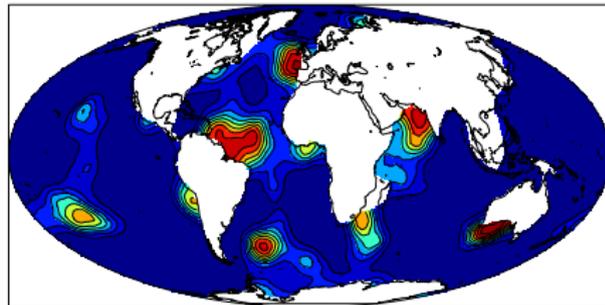
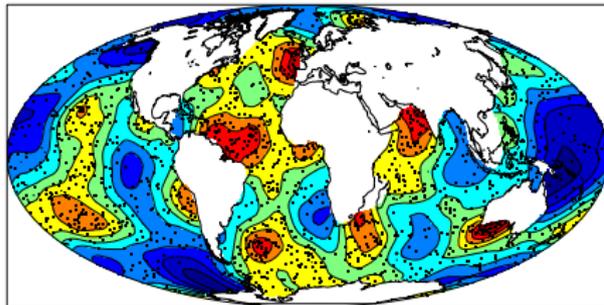
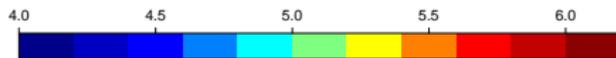
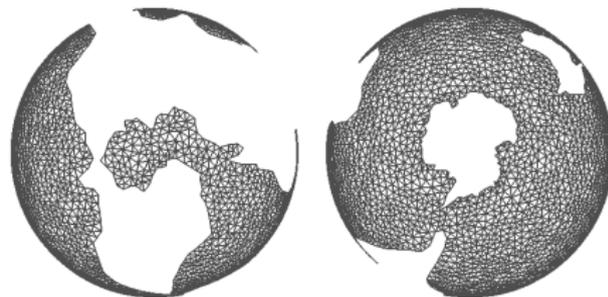
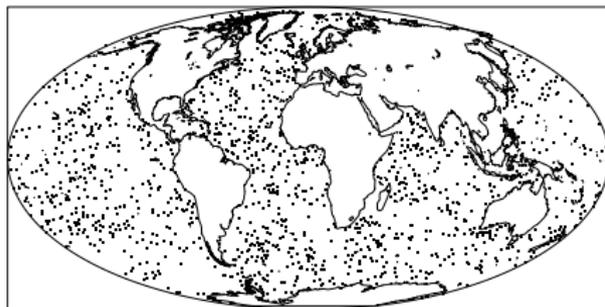
$$\begin{aligned}
 p(\mathbf{u} \mid \boldsymbol{\theta}) &\sim \mathcal{N}(\boldsymbol{\mu}_u, \mathbf{Q}_u^{-1}), & \mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta} &\sim p(\mathbf{y} \mid \mathbf{u}) \\
 p_G(\mathbf{u} \mid \mathbf{y}, \boldsymbol{\theta}) &\sim \mathcal{N}(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{Q}}^{-1}) \\
 \mathbf{0} &= \nabla_{\mathbf{u}} \{ \ln p(\mathbf{u} \mid \boldsymbol{\theta}) + \ln p(\mathbf{y} \mid \mathbf{u}) \} \Big|_{\mathbf{u}=\tilde{\boldsymbol{\mu}}} \\
 \tilde{\mathbf{Q}} &= \mathbf{Q}_u - \nabla_{\mathbf{u}}^2 \ln p(\mathbf{y} \mid \mathbf{u}) \Big|_{\mathbf{u}=\tilde{\boldsymbol{\mu}}}
 \end{aligned}$$

Direct Bayesian inference with INLA (r-inla.org)

$$\begin{aligned}
 \tilde{p}(\boldsymbol{\theta} \mid \mathbf{y}) &\propto \frac{p(\boldsymbol{\theta})p(\mathbf{u} \mid \boldsymbol{\theta})p(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta})}{p_G(\mathbf{u} \mid \mathbf{y}, \boldsymbol{\theta})} \Big|_{\mathbf{u}=\tilde{\boldsymbol{\mu}}} \\
 \tilde{p}(\mathbf{u}_i \mid \mathbf{y}) &\propto \int p_{GG}(\mathbf{u}_i \mid \mathbf{y}, \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\theta} \mid \mathbf{y}) d\boldsymbol{\theta}
 \end{aligned}$$

The main limiting factors for the INLA method are the number of latent variables and the number model parameters (and to a much lesser degree the number of observations). Here we'll concentrate on the latent variables.

SPDE based inference for point process data



A multiscale model example

- ▶ A temporally slow, simplified stochastic heat equation (non-separable)

$$\frac{\partial}{\partial t} z(\mathbf{s}, t) - \gamma_z \nabla \cdot \nabla z(\mathbf{s}, t) = \mathcal{E}(\mathbf{s}, t)$$

$$(1 - \gamma_{\mathcal{E}} \nabla \cdot \nabla) \mathcal{E}(\mathbf{s}, t) = \mathcal{W}_{\mathcal{E}}(\mathbf{s}, t)$$

- ▶ A temporally quick, spatially non-stationary SPDE/GMRF (separable)

$$\left(\frac{\partial}{\partial t} + \gamma_t \right) (\kappa(\mathbf{s})^2 - \nabla \cdot \nabla) (\tau(\mathbf{s}) a(\mathbf{s}, t)) = \mathcal{W}_a(\mathbf{s}, t)$$

- ▶ Measurements

$y_i = a(\mathbf{s}_i, t_i) + z(\mathbf{s}_i, t_i) + \epsilon_i$, discretised into

$$\mathbf{y} = \mathbf{A}(\mathbf{a} + (\mathbf{B} \otimes \mathbf{I})\mathbf{z}) + \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{\epsilon}^{-1})$$

where \mathbf{B} maps from long-term basis functions to short-term, and \mathbf{A} maps from short-term basis functions to the observations.

The posterior precision can be formulated for $(\mathbf{a} + \mathbf{z}, \mathbf{z}) | \mathbf{y}$:

$$\mathbf{Q}_{(\mathbf{a} + \mathbf{z}, \mathbf{z}) | \mathbf{y}} = \begin{bmatrix} \mathbf{Q}_t \otimes \mathbf{Q}_a + \mathbf{A}^{\top} \mathbf{Q}_{\epsilon} \mathbf{A} & -\mathbf{Q}_t \mathbf{B} \otimes \mathbf{Q}_a \\ -\mathbf{B}^{\top} \mathbf{Q}_t \otimes \mathbf{Q}_a & \mathbf{Q}_z + \mathbf{B}^{\top} \mathbf{Q}_t \mathbf{B} \otimes \mathbf{Q}_a \end{bmatrix}$$

Locally isotropic non-stationary precision construction

Finite element construction of basis weight precision

Non-stationary SPDE:

$$(\kappa(\mathbf{s}))^2 - \nabla \cdot \nabla (\tau(\mathbf{s})u(\mathbf{s})) = \mathcal{W}(\mathbf{s})$$

The SPDE parameters are constructed via spatial covariates:

$$\log \tau(\mathbf{s}) = b_0^\tau(\mathbf{s}) + \sum_{j=1}^p b_j^\tau(\mathbf{s})\theta_j, \quad \log \kappa(\mathbf{s}) = b_0^\kappa(\mathbf{s}) + \sum_{j=1}^p b_j^\kappa(\mathbf{s})\theta_j$$

Finite element calculations give

$$\mathbf{T} = \text{diag}(\tau(\mathbf{s}_i)), \quad \mathbf{K} = \text{diag}(\kappa(\mathbf{s}_i))$$

$$C_{ii} = \int \psi_i(\mathbf{s}) d\mathbf{s}, \quad G_{ij} = \int \nabla \psi_i(\mathbf{s}) \cdot \nabla \psi_j(\mathbf{s}) d\mathbf{s}$$

$$\mathbf{Q} = \mathbf{T} (\mathbf{K}^2 \mathbf{C} \mathbf{K}^2 + \mathbf{K}^2 \mathbf{G} + \mathbf{G} \mathbf{K}^2 + \mathbf{G} \mathbf{C}^{-1} \mathbf{G}) \mathbf{T}$$

Combining this with an AR(1) discretisation of the temporal operator, we get $\mathbf{Q}_t \otimes \mathbf{Q}_a$. More on \mathbf{Q}_z tomorrow.

Posterior calculations

Write $x = (a + z, z)$ for the full latent field.

$$Q_{x|y} = \begin{bmatrix} Q_t \otimes Q_a + A^\top Q_\epsilon A & -Q_t B \otimes Q_a \\ -B^\top Q_t \otimes Q_a & Q_z + B^\top Q_t B \otimes Q_a \end{bmatrix}$$

can be pseudo-Cholesky-factorised:

$$Q_{x|y} = \tilde{L}_{x|y} \tilde{L}_{x|y}^\top, \quad \tilde{L}_{x|y} = \begin{bmatrix} L_t \otimes L_a & \mathbf{0} & A^\top L_\epsilon \\ -B^\top L_t \otimes L_a & \tilde{L}_z & \mathbf{0} \end{bmatrix}$$

Posterior expectation, samples, and marginal variances (with $\tilde{A} = [A \quad \mathbf{0}]$):

$$Q_{x|y}(\mu_{x|y} - \mu_x) = \tilde{A}^\top Q_\epsilon (y - \tilde{A}\mu_x),$$

$$Q_{x|y}(x - \mu_{x|y}) = \tilde{L}_{x|y} w, \quad w \sim \mathcal{N}(\mathbf{0}, I), \quad \text{or}$$

$$Q_{x|y}(x - \mu_x) = \tilde{A}^\top Q_\epsilon (y - \tilde{A}\mu_x) + \tilde{L}_{x|y} w, \quad w \sim \mathcal{N}(\mathbf{0}, I),$$

$$\text{Var}(x_i|y) = \text{diag}(\text{inla.qinv}(Q_{x|y})) \quad (\text{requires Cholesky})$$

Preconditioning for e.g. conjugate gradient solutions

Solving $Qx = b$ is equivalent to solving $M^{-1}Qx = M^{-1}b$. Choosing M^{-1} as an approximate inverse to Q gives a less ill-conditioned system. Only the *action* of M^{-1} is needed, e.g. one or more fixed point iterations:

Block Jacobi and Gauss-Seidel preconditioning

$$\text{Matrix split: } Q_{x|y} = L + D + L^\top$$

$$\text{Jacobi: } x^{(k+1)} = D^{-1} \left(-(L + L^\top)x^{(k)} + b \right)$$

$$\text{Gauss-Seidel: } x^{(k+1)} = (L + D)^{-1} \left(-L^\top x^{(k)} + b \right)$$

Remark: Block Gibbs sampling for a GMRF posterior

With $Q = Q_{x|y}$, $b = A^\top Q_\epsilon (y - A\mu_x)$ and $\tilde{x} = x - \mu_x$,

$$\tilde{x}^{(k+1)} = (L + D)^{-1} \left(-L^\top \tilde{x}^{(k)} + b + \tilde{L}_D w \right), \quad w \sim \mathcal{N}(0, I)$$

Gauss-Seidel and Gibbs are both inefficient on their own, but G-S leads to useful preconditioners. Convergence testing is much easier for linear solvers than for MCMC.

Multiscale Schur complement approximation

Solving $Q_{x|y}x = b$ can be formulated using two solves with the upper block $Q_t \otimes Q_a + A^\top Q_\epsilon A$, and one solve with the *Schur complement*

$$Q_z + B^\top Q_t B \otimes Q_a - B^\top Q_t \otimes Q_a \left(Q_t \otimes Q_a + A^\top Q_\epsilon A \right)^{-1} Q_t B \otimes Q_a$$

By mapping the fine scale model onto the coarse basis used for the coarse model, we get an *approximate* (and sparse) Schur solve via

$$\begin{bmatrix} \tilde{Q}_B + \tilde{B}^\top A^\top Q_\epsilon A \tilde{B} & -\tilde{Q}_B \\ -\tilde{Q}_B & Q_z + \tilde{Q}_B \end{bmatrix} \begin{bmatrix} \text{ignored} \\ z \end{bmatrix} = \begin{bmatrix} 0 \\ \tilde{b} \end{bmatrix}$$

where $\tilde{B} = B \otimes I$, $\tilde{Q}_B = B^\top Q_t B \otimes Q_a$, and the block matrix can be interpreted as the precision of a bivariate field on a common, coarse spatio-temporal scale.

Each block may be very large even on the coarse scale, and we still need a way to solve with the fine scale upper block!

Multigrid

Construct a sequence of increasingly detailed models,

$$\left(Q^{(0)}, Q^{(1)}, \dots, Q^{(L)} \right).$$

Basic idea:

- ▶ On each level, a simple local fixed point iteration can eliminate small scale residual errors efficiently, but not large scale errors.
- ▶ Project the residual onto the next coarse level, where the large scale is now small, and then interpolate the result back onto the finer level.
- ▶ On the coarsest level, solve the exact problem.

Simple multigrid model traversal: $L = 4, 3, 2, 1, 0, 1, 2, 3, 4 = L$

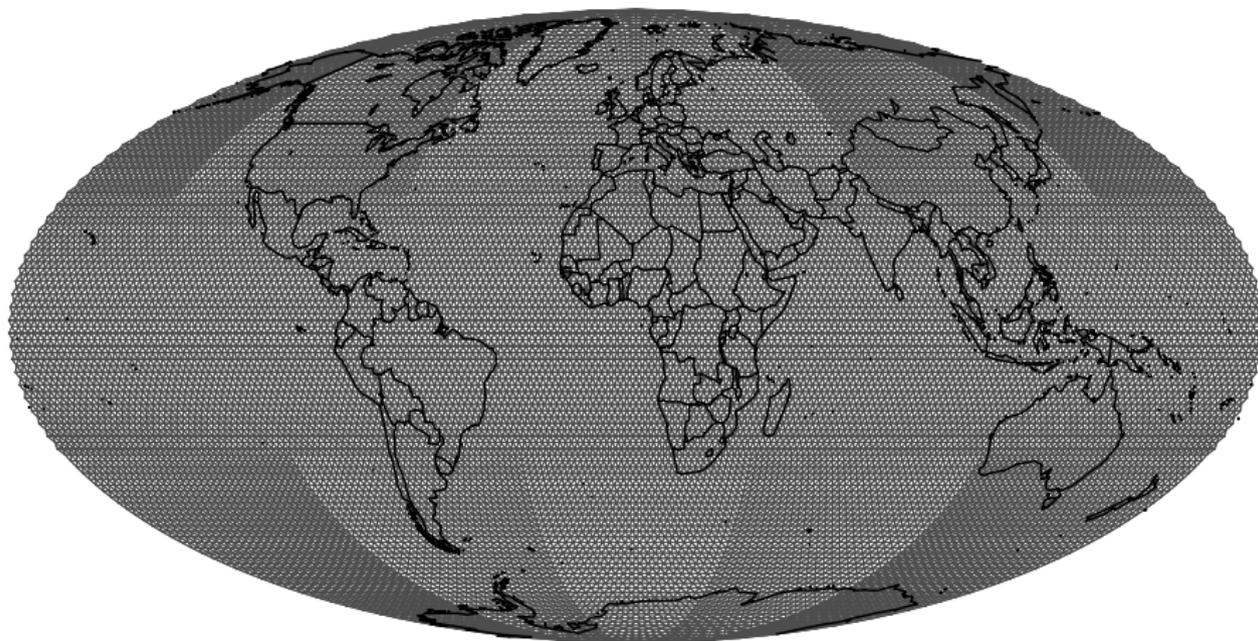
Full multigrid: $L = 4, 3, 2, 1, 0, 1, 0, 1, 2, 1, 0, 1, 2, 3, 2, 1, 0, 1, 2, 3, 4 = L$

In theory, full multigrid can be $\mathcal{O}(n)$, but the constant may be very large, and toy model testing indicates that the high operator order is a problem.

Can be used as complete solver with small tolerance, or as preconditioner with large tolerance.

Finite element mesh

Triangulation mesh



Finite element mesh (spot the R map plotting surprise!)



Domain decomposition

- ▶ Divide the domain into a collection of overlapping subdomain blocks
- ▶ Solve a local problem, e.g. the conditional solution, maintaining coherence by enforcing constraints on overlapping nodes.

Monte Carlo variance reduction for posterior variances

$$E(\mathbf{x}_i | \mathbf{y}) = E(E(\mathbf{x}_i | \mathbf{y}, \mathbf{x}_{\not\subseteq \text{subblock}}))$$

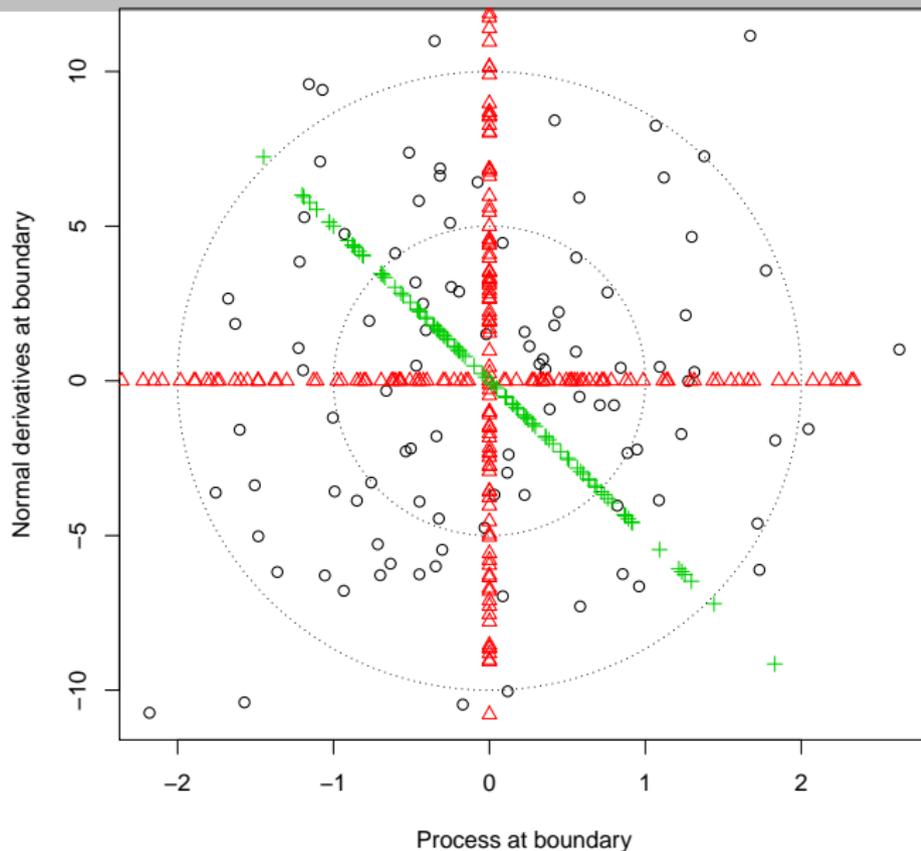
$$\text{Var}(\mathbf{x}_i | \mathbf{y}) = \text{Var}(\mathbf{x}_i | \mathbf{y}, \mathbf{x}_{\not\subseteq \text{subblock}}) + \text{Var}(E(\mathbf{x}_i | \mathbf{y}, \mathbf{x}_{\not\subseteq \text{subblock}}))$$

Also works for linear combinations, with some complications

Subdomain boundary adjustment (new idea)

- ▶ Apply *stochastic boundary correction* for each subdomain
- ▶ Solve the full local problem, reusing the appropriate randomness for overlapping subdomains
- ▶ Blend the results for overlapping domains.
- ▶ Apply this as a preconditioner in an iterative solver

All deterministic boundary conditions are 'inappropriate'



Stationary stochastic boundary adjustment

Recall the Matérn generating SPDE

$$(\kappa^2 - \nabla \cdot \nabla)^{\alpha/2} u(\mathbf{s}) = \mathcal{W}(\mathbf{s})$$

RKHS inner product for Matérn precisions on \mathbb{R}^d :

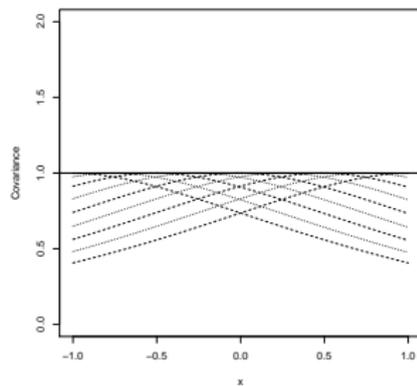
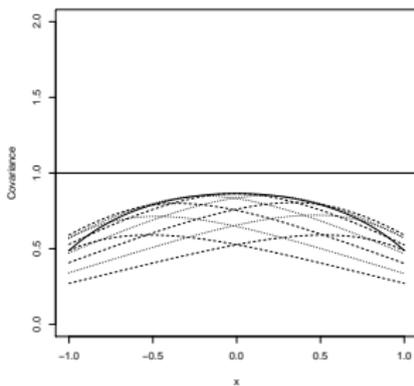
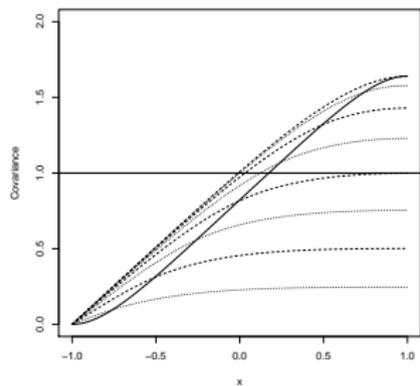
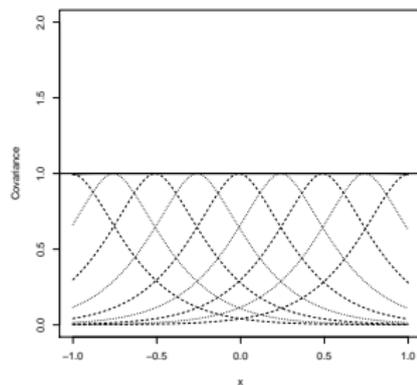
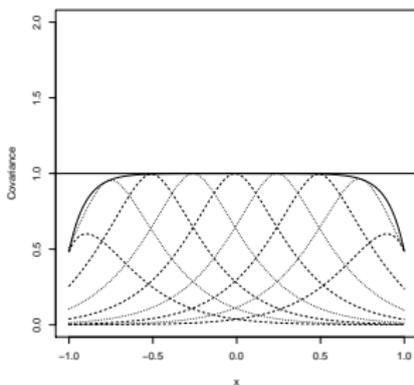
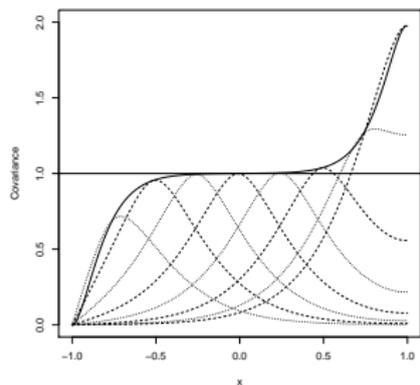
$$\langle f, g \rangle_{H(\Omega)} = \sum_{k=0}^{\alpha} \binom{\alpha}{k} \kappa^{2\alpha-2k} \langle \nabla^k f, \nabla^k g \rangle_{\Omega}$$

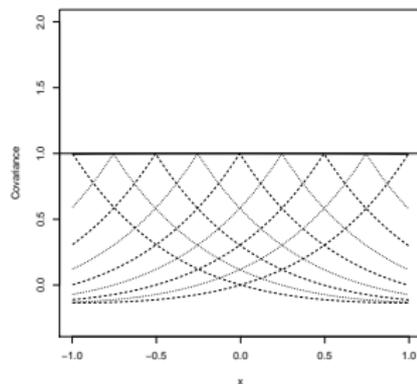
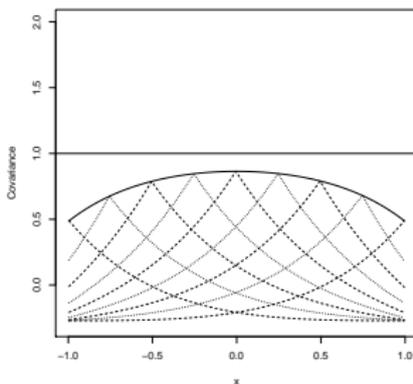
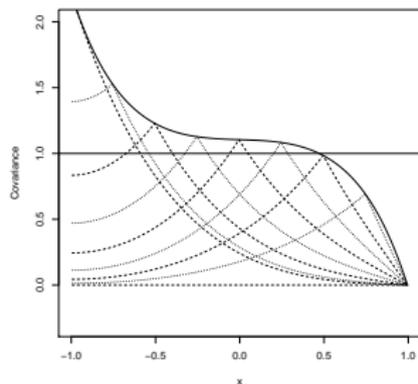
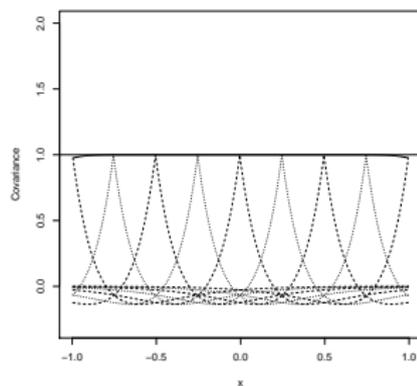
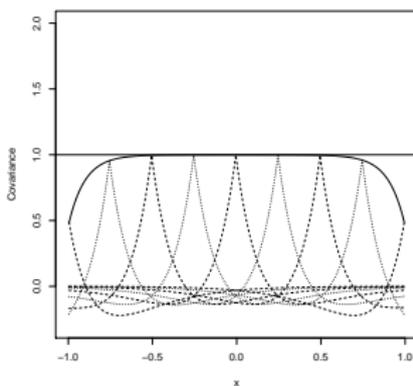
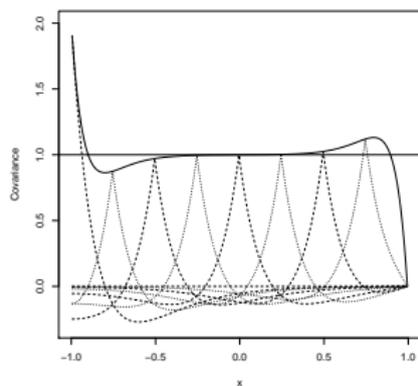
Boundary adjusted precision operator on a compact subdomain, where \mathcal{P} projects onto the operator null-space:

$$\begin{aligned} \mathcal{Q}_{\Omega}(f, g) &= \langle f, g \rangle_{H(\Omega)} - \langle \mathcal{P}f, \mathcal{P}g \rangle_{H(\Omega)} + \mathcal{Q}_{\mathcal{P};\partial\Omega}(\mathcal{P}f, \mathcal{P}g) \\ &= \langle f - \mathcal{P}f, g - \mathcal{P}g \rangle_{H(\Omega)} + \mathcal{Q}_{\mathcal{P};\partial\Omega}(\mathcal{P}f, \mathcal{P}g) \end{aligned}$$

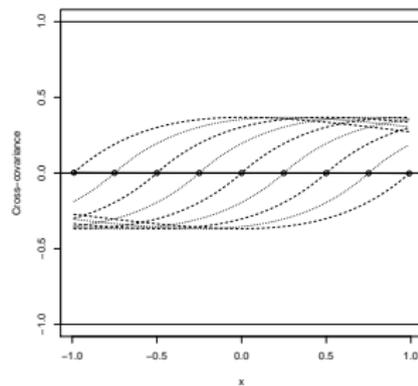
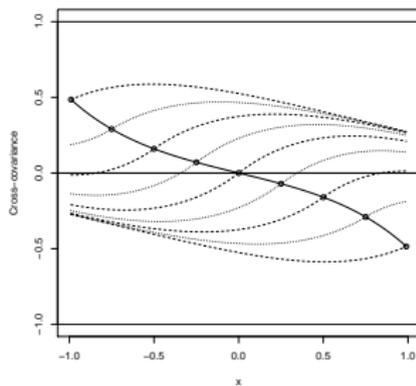
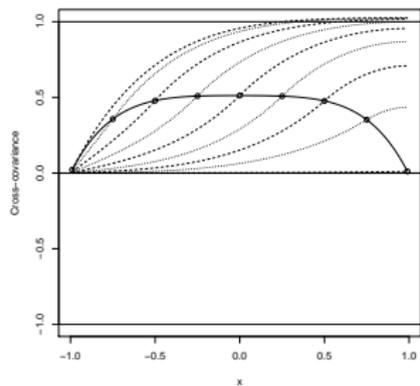
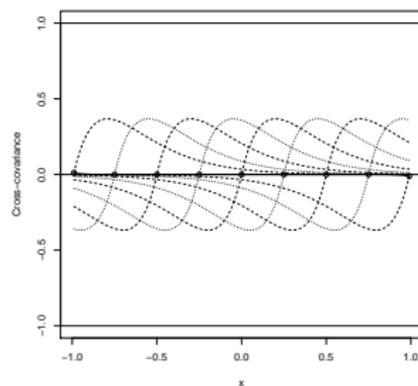
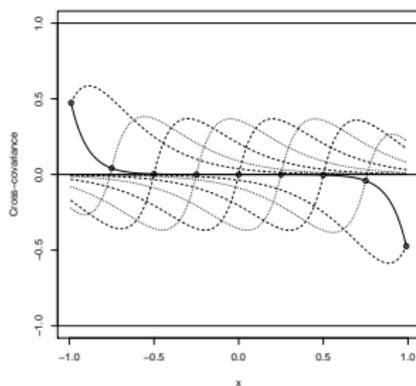
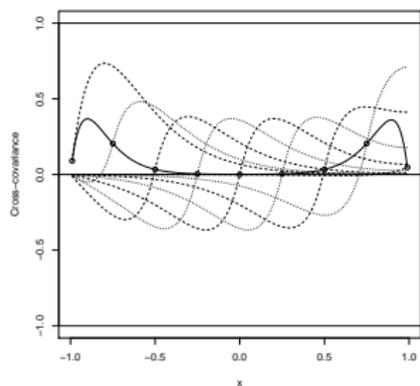
Note that $\mathcal{Q}_{\mathcal{P};\partial\Omega}(\mathcal{P}f, \mathcal{P}g)$ may involve normal derivatives at the boundary.

Covariances (D&N, Robin, Stoch) for $\kappa = 5$ and 1

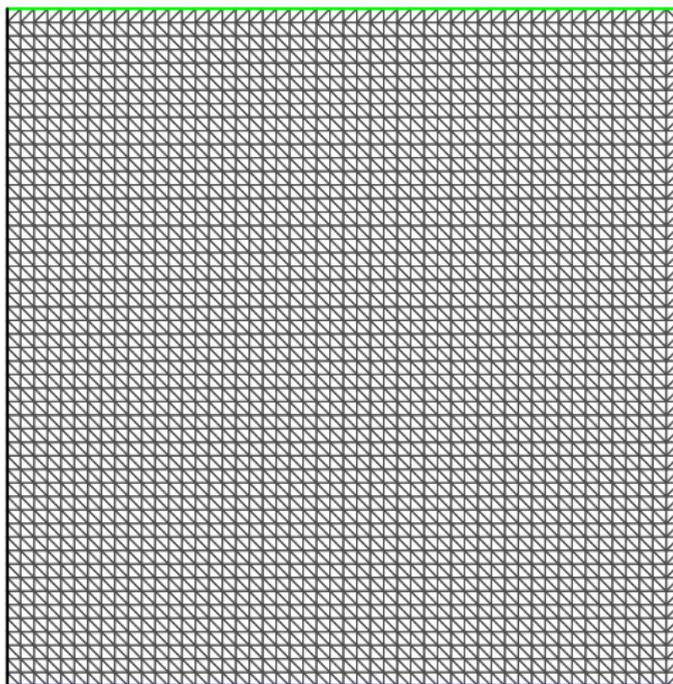


Derivative covariances (D&N, Robin, Stoch) for $\kappa = 5$ and 1

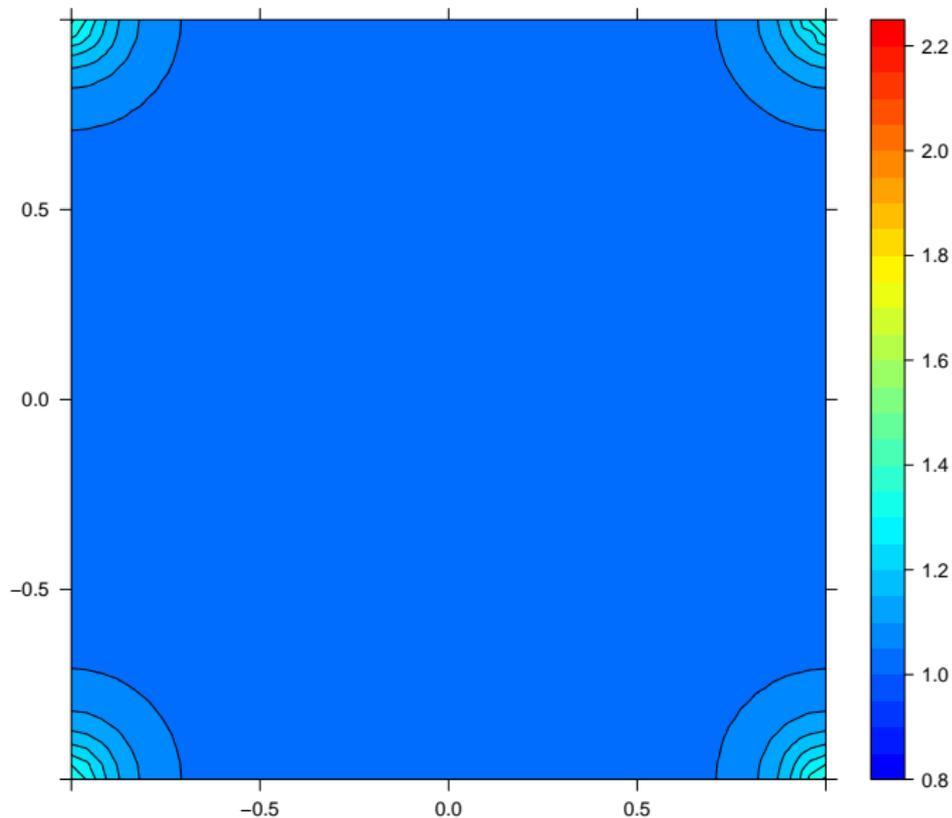
Process-derivative cross-covariances (D&N, Robin, Stoch)



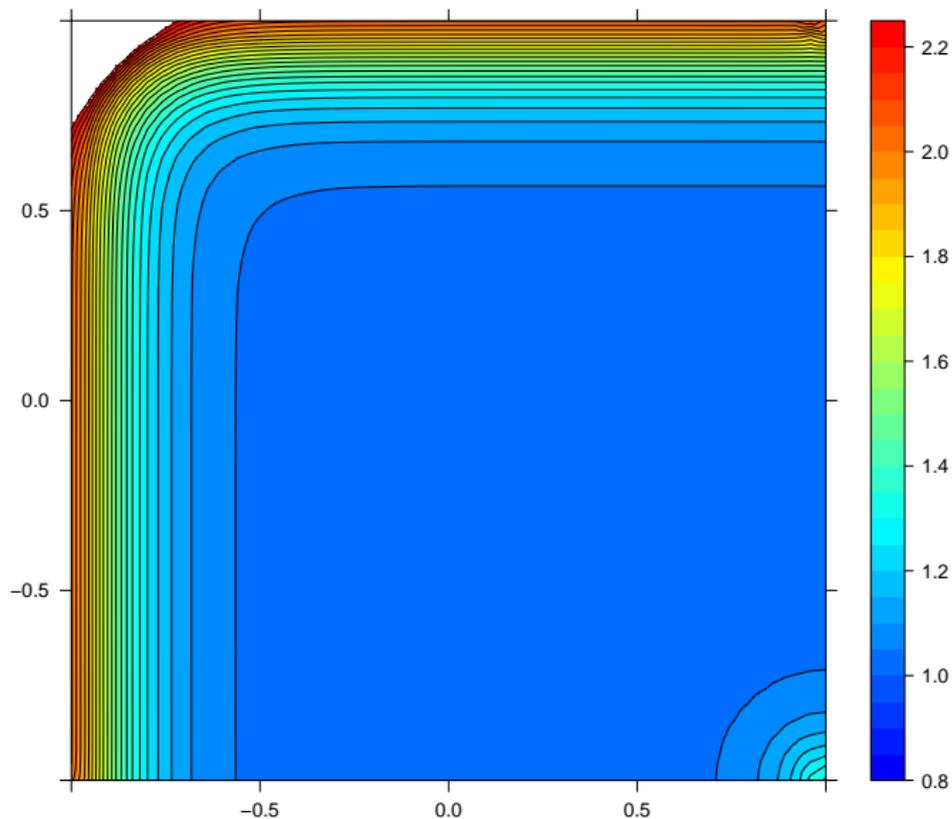
Square domain, basis triangulation



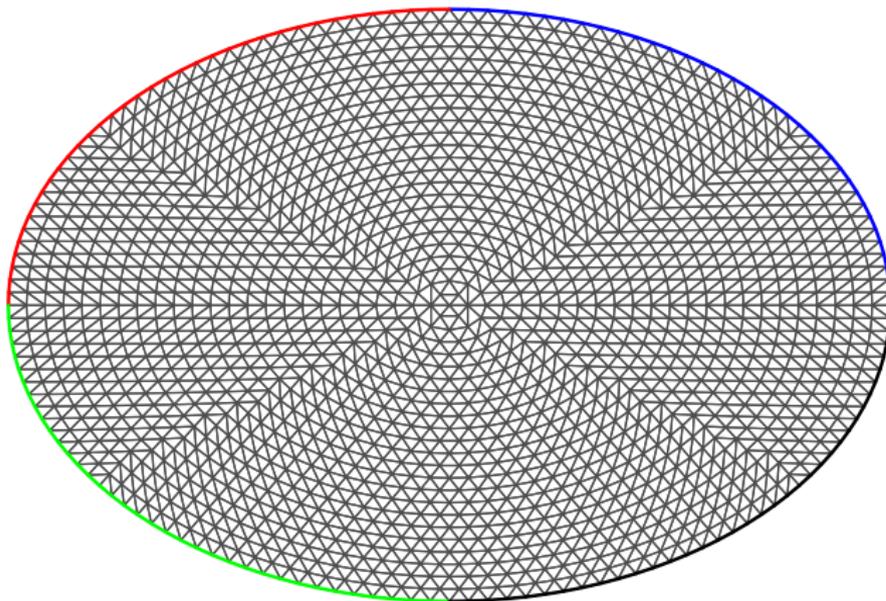
Square domain, stochastic boundary (variances)



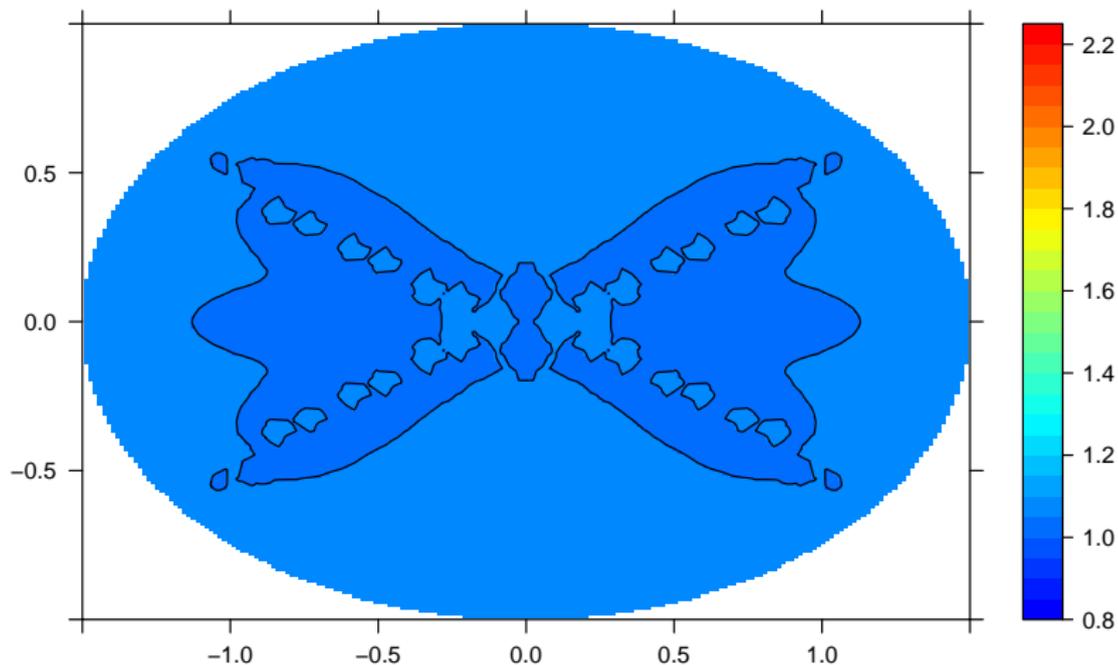
Square domain, mixed boundary (variances)



Elliptical domain, basis triangulation



Elliptical domain, stochastic boundary (variances)



References

References (see also r-inla.org)

- ▶ F. Lindgren, H. Rue and J. Lindström (2011), *An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach (with discussion)*, JRSSB, 73(4), 423–498.
- ▶ Preliminary draft book chapter on GMRF computation basics: <http://people.bath.ac.uk/f1353/tmp/gmrf.pdf>
- ▶ R. Ingebrigtsen, F. Lindgren, I. Steinsland (2013), *Spatial models with explanatory variables in the dependence structure*, Spatial Statistics, In Press (available online).
- ▶ G-A. Fuglstad, F. Lindgren, D. Simpson, H. Rue (2013), *Exploring a new class of non-stationary spatial Gaussian random fields with varying local anisotropy*, accepted for Statistica Sinica, arXiv:1304.6949
- ▶ D. Bolin, F. Lindgren (2013), *Excursion and contour uncertainty regions for latent Gaussian models*, JRSSB, in press. Accepted version at arXiv:1211.3946.
CRAN package: `excursions`