Large scale spatial statistics with stochastic PDEs

Finn Lindgren

Håvard Rue, Johan Lindström, Daniel Simpson, Janine Illian, David Bolin, Michela Cameletti, Geir-Arne Fuglstad, Rikke Ingebrigtsen, Xiangping Hu, Peter Guttorp

UNIVERSITY OF BATH

GRF-Sim, Bern, 26 November 2014

“Big” data

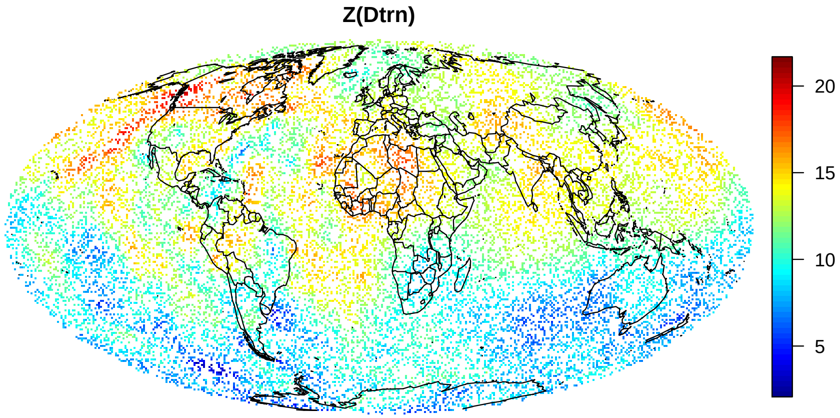

Synthetic data mimicking satellite based CO_2 measurements.
Iregular data locations, uneven coverage, and all scales need to be handled.

Sparse spatial coverage of temperature measurements

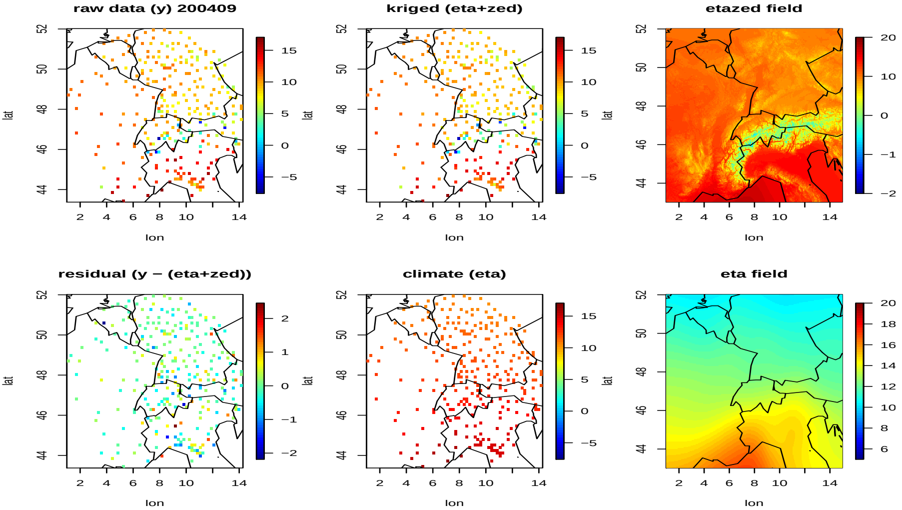

Regional observations: ≈ 20,000,000 from daily timeseries over 160 years

Note: This is a small *subset* of the full data!

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 $u(\boldsymbol{s})$, $\boldsymbol{s} \in \Omega$, or $u(\boldsymbol{s}, t)$, $(\boldsymbol{s}, t) \in \Omega \times \mathbb{T}$.
- Observations y_i. In the simplest setting, $y_i = u(\boldsymbol{s}_i) + \epsilon_i$, but more generally $y_i \sim$ GLMM, with $u(\cdot)$ as a structured random effect.
- I'll restrict this talk to Gaussian latent models.
- To simplify the presentation, most of the measurement models are simplified to simply iid measurement errors, but keep in mind that we may eventually need methods that can handle individual long-term random effects for each weather station/buoy/ship/satellite.

Covariance functions and stochastic PDEs

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

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

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

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

$\mathcal{W}(\cdot)$ 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 and discrete Markov properties

Markov properties

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

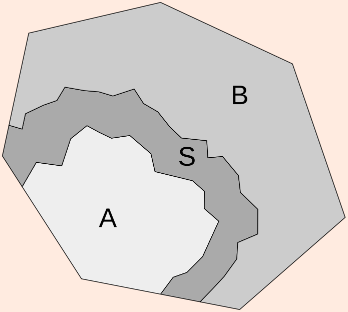

Solutions to
$$\left(\kappa^2 - \nabla \cdot \nabla\right)^{\alpha/2} u(\boldsymbol{s}) = \mathcal{W}(\boldsymbol{s})$$
are Markov when α is an integer.
(Rozanov, 1977)

Discrete representations ($\boldsymbol{Q} = \boldsymbol{\Sigma}^{-1}$):
$$\boldsymbol{Q}_{AB} = \boldsymbol{0}$$
$$\boldsymbol{Q}_{A|S,B} = \boldsymbol{Q}_{AA}$$
$$\boldsymbol{\mu}_{A|S,B} = \boldsymbol{\mu}_A - \boldsymbol{Q}_{AA}^{-1} \boldsymbol{Q}_{AS}(\boldsymbol{u}_S - \boldsymbol{\mu}_S)$$

Continuous domain Markov approximations

Continuous Markovian spatial models (Lindgren et al, 2011)

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

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

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

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

Conditional distribution in a Gaussian model

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

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

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

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

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)

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

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

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

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)

Non-stationary field

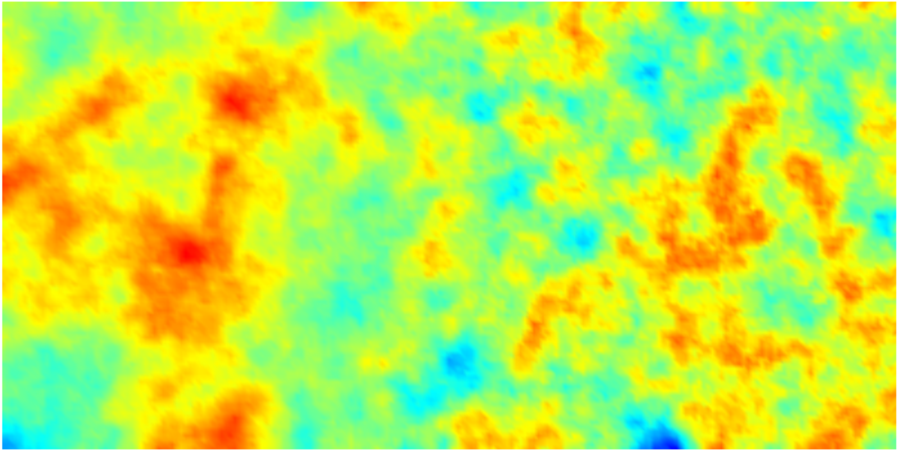

$$(\kappa(\boldsymbol{s})^2 - \nabla \cdot \nabla) u(\boldsymbol{s}) = \kappa(\boldsymbol{s}) \mathcal{W}(\boldsymbol{s}), \quad \boldsymbol{s} \in \Omega$$

Anisotropic field on a globe via vector parameter field

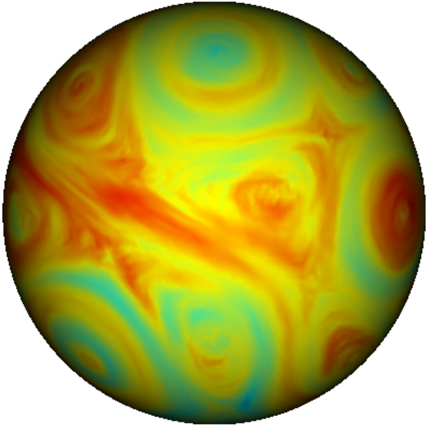

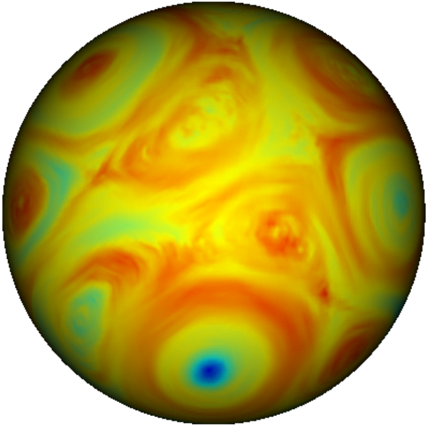

$$(\kappa(\boldsymbol{s})^2 - \nabla \cdot \boldsymbol{H}(\boldsymbol{s})\nabla)u(\boldsymbol{s}) = \kappa(\boldsymbol{s})\mathcal{W}(\boldsymbol{s}), \quad \boldsymbol{s} \in \Omega$$

Covariances for four reference points

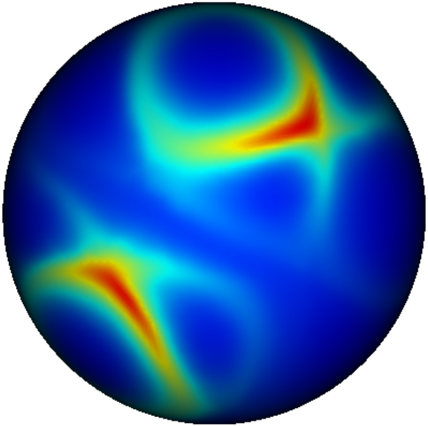

Climate and weather model (simplified)

- Climate process, simplified stochastic heat equation

$$
\begin{aligned}
\frac{\partial}{\partial t} z(\boldsymbol{s}, t) - \nabla \cdot \nabla z(\boldsymbol{s}, t) &= \mathcal{E}(\boldsymbol{s}, t) \\
(1 - \gamma_{\mathcal{E}} \nabla \cdot \nabla) \mathcal{E}(\boldsymbol{s}, t) &= \mathcal{W}_{\mathcal{E}}(\boldsymbol{s}, t)
\end{aligned}
$$

- Weather anomaly, non-stationary spatial SPDE/GMRF
  $(\kappa(\boldsymbol{s})^2 - \nabla \cdot \nabla) (\tau(\boldsymbol{s}) a(\boldsymbol{s}, t)) = \mathcal{W}_a(\boldsymbol{s}, t)$
- Temperature measurements from one or several sources
  $\boldsymbol{y} = \boldsymbol{A}(\boldsymbol{a} + (\boldsymbol{B} \otimes \boldsymbol{I})\boldsymbol{z}) + \boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_{\epsilon}^{-1})$

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

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

Locally isotropic non-stationary precision construction

Finite element construction of basis weight precision

Non-stationary SPDE:

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

The SPDE parameters are constructed via spatial covariates:

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

Finite element calculations give

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

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

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

For the temporally independent anomalies, we get $\boldsymbol{I}\otimes\boldsymbol{Q}_a$

GMRF precision for simplified stochastic heat equation

$$
\begin{aligned}
\boldsymbol{Q}_z &= \boldsymbol{M}_2^{(t)} \otimes \boldsymbol{M}_0^{(s)} + \boldsymbol{M}_1^{(t)} \otimes \boldsymbol{M}_1^{(s)} + \boldsymbol{M}_0^{(t)} \otimes \boldsymbol{M}_2^{(s)} \\
\boldsymbol{M}_0^{(s)} &= \boldsymbol{C} + \gamma_{\mathcal{E}} \boldsymbol{G} \\
\boldsymbol{M}_1^{(s)} &= \boldsymbol{G} + \gamma_{\mathcal{E}} \boldsymbol{G} \boldsymbol{C}^{-1} \boldsymbol{G} \\
\boldsymbol{M}_2^{(s)} &= \boldsymbol{G} \boldsymbol{C}^{-1} \boldsymbol{G} + \gamma_{\mathcal{E}} \boldsymbol{G} \boldsymbol{C}^{-1} \boldsymbol{G} \boldsymbol{C}^{-1} \boldsymbol{G}
\end{aligned}
$$

Ignoring the degenerate aspect of the model, the precision structure can be used to formulate sampling as (remember the first talk of the workshop!)

$$
\boldsymbol{Q}_z \boldsymbol{z} = \widetilde{\boldsymbol{L}}_z \boldsymbol{w}, \quad \boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})
$$

where $\widetilde{\boldsymbol{L}}_z$ is a pseudo Cholesky factor,

$$
\begin{aligned}
\widetilde{\boldsymbol{L}}_z = \Big[& \left[\boldsymbol{L}_2^{(t)} \otimes \boldsymbol{L}_C, \quad \boldsymbol{L}_1^{(t)} \otimes \boldsymbol{L}_G, \quad \boldsymbol{L}_0^{(t)} \otimes \boldsymbol{G} \boldsymbol{L}_C^{-\top} \right], \\
& \gamma_{\mathcal{E}}^{1/2} \left[\boldsymbol{L}_2^{(t)} \otimes \boldsymbol{L}_G, \quad \boldsymbol{L}_1^{(t)} \otimes \boldsymbol{G} \boldsymbol{L}_C^{-\top}, \quad \boldsymbol{L}_0^{(t)} \otimes \boldsymbol{G} \boldsymbol{C}^{-1} \boldsymbol{L}_G \right] \Big]
\end{aligned}
$$

Posterior calculations

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

$$\boldsymbol{Q}_{x|y} = \begin{bmatrix} \boldsymbol{I} \otimes \boldsymbol{Q}_a + \boldsymbol{A}^\top \boldsymbol{Q}_\epsilon \boldsymbol{A} & -\boldsymbol{B} \otimes \boldsymbol{Q}_a \\ -\boldsymbol{B}^\top \otimes \boldsymbol{Q}_a & \boldsymbol{Q}_z + \boldsymbol{B}^\top \boldsymbol{B} \otimes \boldsymbol{Q}_a \end{bmatrix}$$

can be pseudo-Cholesky-factorised:

$$\boldsymbol{Q}_{x|y} = \widetilde{\boldsymbol{L}}_{x|y} \widetilde{\boldsymbol{L}}_{x|y}^\top, \qquad \widetilde{\boldsymbol{L}}_{x|y} = \begin{bmatrix} \boldsymbol{I} \otimes \boldsymbol{L}_a & \boldsymbol{0} & \boldsymbol{A}^\top \boldsymbol{L}_\epsilon \\ -\boldsymbol{B} \otimes \boldsymbol{L}_a & \widetilde{\boldsymbol{L}}_z & \boldsymbol{0} \end{bmatrix}$$

Posterior expectation, samples, and marginal variances:

$$\boldsymbol{Q}_{x|y}(\boldsymbol{\mu}_{x|y} - \boldsymbol{\mu}_x) = \boldsymbol{A}^\top \boldsymbol{Q}_\epsilon (\boldsymbol{y} - \boldsymbol{\mu}_x),$$

$$\boldsymbol{Q}_{x|y}(\boldsymbol{x} - \boldsymbol{\mu}_{x|y}) = \widetilde{\boldsymbol{L}}_{x|y} \boldsymbol{w}, \quad \boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}), \quad \text{or}$$

$$\boldsymbol{Q}_{x|y}(\boldsymbol{x} - \boldsymbol{\mu}_x) = \boldsymbol{A}^\top \boldsymbol{Q}_\epsilon (\boldsymbol{y} - \boldsymbol{\mu}_x) + \widetilde{\boldsymbol{L}}_{x|y} \boldsymbol{w}, \quad \boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}),$$

$$\mathsf{Var}(x_i | \boldsymbol{y}) = \texttt{diag(inla.qinv(}\boldsymbol{Q}_{x|y}\texttt{))} \quad \text{(requires Cholesky)}$$

Preconditioning for e.g. conjugate gradient solutions

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

Block Jacobi and Gauss-Seidel preconditioning

Matrix split: $\boldsymbol{Q}_{x|y} = \boldsymbol{L} + \boldsymbol{D} + \boldsymbol{L}^\top$

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

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

Remark: Block Gibbs sampling for a GMRF posterior

With $\boldsymbol{Q} = \boldsymbol{Q}_{x|y}$, $\boldsymbol{b} = \boldsymbol{A}^\top\boldsymbol{Q}_\epsilon(\boldsymbol{y} - \boldsymbol{\mu}_x)$ and $\widetilde{\boldsymbol{x}} = \boldsymbol{x} - \boldsymbol{\mu}_x$,

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

Multiscale Schur complement approximation

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

$$\boldsymbol{Q}_z + \boldsymbol{B}^\top \boldsymbol{B} \otimes \boldsymbol{Q}_a - \boldsymbol{B}^\top \otimes \boldsymbol{Q}_a \left(\boldsymbol{I} \otimes \boldsymbol{Q}_a + \boldsymbol{A}^\top \boldsymbol{Q}_\epsilon \boldsymbol{A}\right)^{-1} \boldsymbol{B} \otimes \boldsymbol{Q}_a$$

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

$$\begin{bmatrix} \widetilde{\boldsymbol{Q}}_B + \widetilde{\boldsymbol{B}}^\top \boldsymbol{A}^\top \boldsymbol{Q}_\epsilon \boldsymbol{A} \widetilde{\boldsymbol{B}} & -\widetilde{\boldsymbol{Q}}_B \\ -\widetilde{\boldsymbol{Q}}_B & \boldsymbol{Q}_z + \widetilde{\boldsymbol{Q}}_B \end{bmatrix} \begin{bmatrix} \text{ignored} \\ \boldsymbol{z} \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \widetilde{\boldsymbol{b}} \end{bmatrix}$$

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

Multigrid

Construct a sequence of increasingly detailed models, $\left(\boldsymbol{Q}^{(0)}, \boldsymbol{Q}^{(1)}, \ldots, \boldsymbol{Q}^{(L)}\right)$.

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 $\mathcal{O}(n)$!

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

Finite element mesh

Triangulation mesh

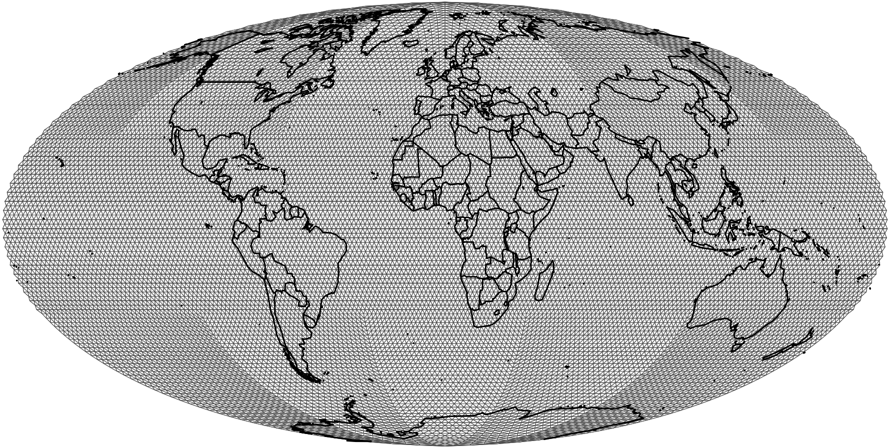

Finite element mesh

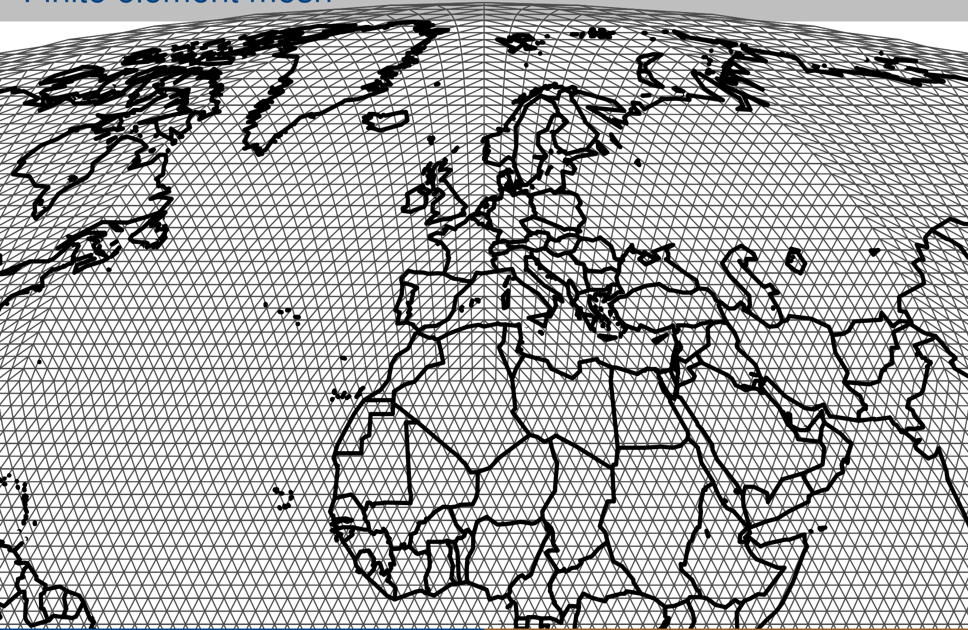

Domain decomposition

- 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

$$\text{Var}(\boldsymbol{x}_i \mid \boldsymbol{y}) = \text{Var}\left(\boldsymbol{x}_i \mid \boldsymbol{y}, \boldsymbol{x}_{\notin\text{subblock}}\right) + \text{Var}\left(\text{E}(\boldsymbol{x}_i \mid \boldsymbol{y}, \boldsymbol{x}_{\notin\text{subblock}})\right)$$

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.

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

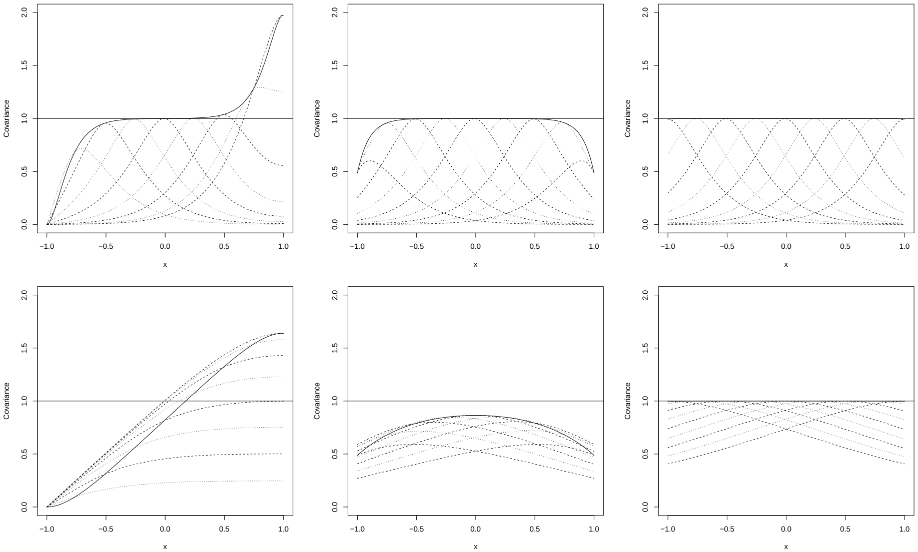

Elliptical domain, basis triangulation

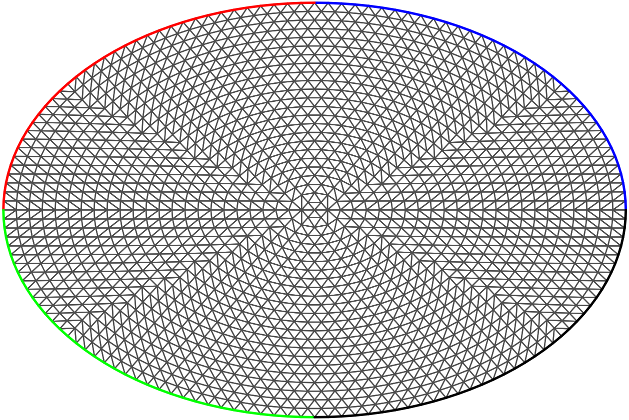

Elliptical domain, stochastic boundary (variances)

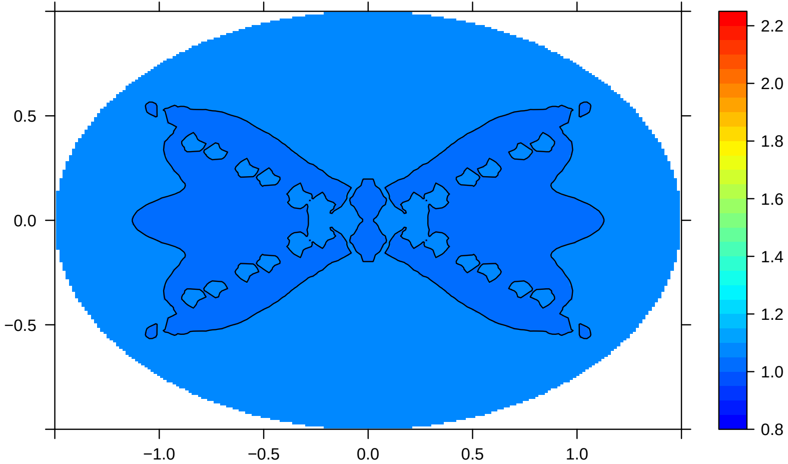

Elliptical domain, mixed boundary (variances)

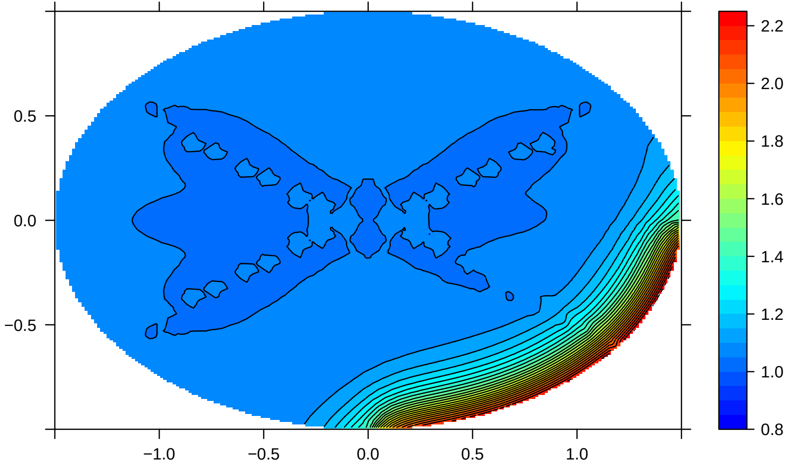

Stationary stochastic boundary adjustment

Recall the Matérn generating SPDE

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

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

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

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

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

Laplace approximations for non-Gaussian observations

Quadratic posterior log-likelihood approximation

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

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

$$
\widetilde{p}(\boldsymbol{\theta} \mid \boldsymbol{y}) \propto \left. \frac{p(\boldsymbol{\theta}) p(\boldsymbol{u} \mid \boldsymbol{\theta}) p(\boldsymbol{y} \mid \boldsymbol{u}, \boldsymbol{\theta})}{p_G(\boldsymbol{u} \mid \boldsymbol{y}, \boldsymbol{\theta})} \right|_{\boldsymbol{u}=\widetilde{\boldsymbol{\mu}}}
$$

$$
\widetilde{p}(\boldsymbol{u}_i \mid \boldsymbol{y}) \propto \int p_{GG}(\boldsymbol{u}_i \mid \boldsymbol{y}, \boldsymbol{\theta}) \widetilde{p}(\boldsymbol{\theta} \mid \boldsymbol{y}) \, d\boldsymbol{\theta}
$$

The latent Gaussian parts to some degree do scale to large non direct methods, but evaluating likelihoods becomes a very challenging problem.

SPDE based inference for point process data

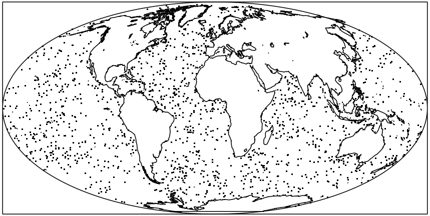

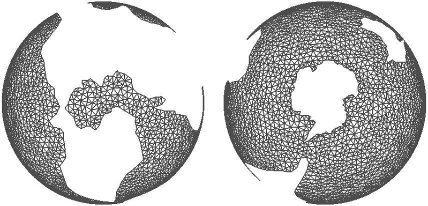

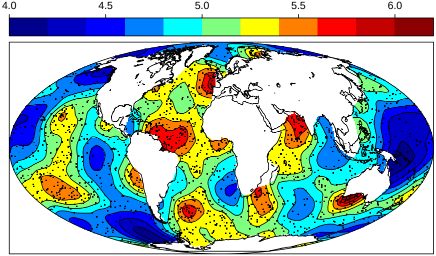

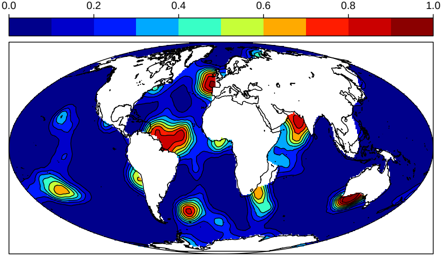

Excursion sets for random fields

Excursion sets

Let $x(s)$, $s \in \Omega$ be a random process. The positive and negative level u excursion sets with probability $1-\alpha$ are

$$E^{+}_{u,\alpha}(x) = \underset{D}{\operatorname{argmax}}\{|D| : \mathsf{Pr}(D \subseteq A^{+}_{u}(x)) \geq 1-\alpha\}.$$
$$E^{-}_{u,\alpha}(x) = \underset{D}{\operatorname{argmax}}\{|D| : \mathsf{Pr}(D \subseteq A^{-}_{u}(x)) \geq 1-\alpha\}.$$

Excursion functions

The positive and negative u excursion functions are given by

$$F^{+}_{u}(s) = \sup\{1-\alpha; s \in E^{+}_{u,\alpha}\},$$
$$F^{-}_{u}(s) = \sup\{1-\alpha; s \in E^{-}_{u,\alpha}\}.$$

PM_{10} exceedances in Piemonte, January 30, 2006

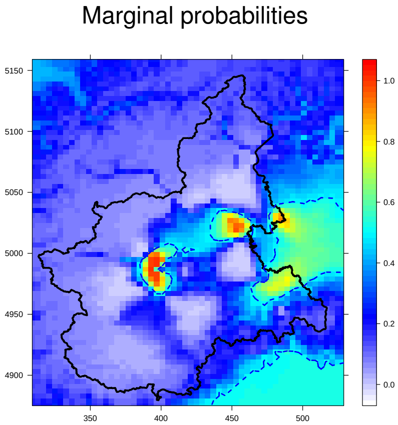

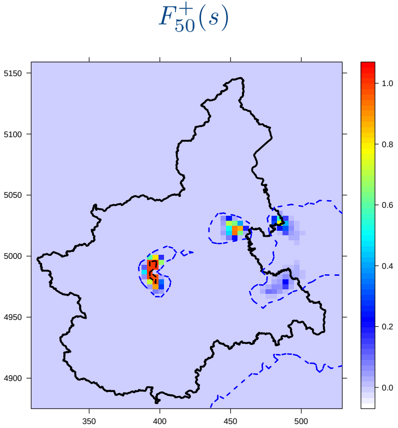

Model estimated with INLA, result passed onward to `excursions()`, evaluating high dimensional GMRF probabilities and finding credible regions. Development version has more user friendly options for continuous domain interpretations.

References

- Rue, H. and Held, L.: Gaussian Markov Random Fields; Theory and Applications; *Chapman & Hall/CRC*, 2005
- Lindgren, F.: Computation fundamentals of discrete GMRF representations of continuous domain spatial models; preliminary book chapter manuscript, 2014,
  `http://people.bath.ac.uk/fl353/tmp/gmrf.pdf`
- Lindgren, F., Rue, H., and Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach (with discussion); *JRSS Series B*, 2011
  Non-CRAN package: `R-INLA` at `http://r-inla.org/`
- Bolin, D. and Lindgren, F.: Excursion and contour uncertainty regions for latent Gaussian models; *JRSS Series B*, 2014, in press. Accepted version at arXiv:1211.3946 and on journal web page.
  CRAN package: `excursions`
  Development: `http://bitbucket.org/davidbolin/excursions`