Non-separable diffusion-based spatio-temporal Gaussian fields

Finn Lindgren (finn.lindgren@ed.ac.uk)

with Elias Krainski, David Bolin, Haakon Bakka, Haavard Rue

THE UNIVERSITY of EDINBURGH

CMStatistics, London, 2022-12-17

The eternal quest for spatial dependence models

- Gaussian random field: $u(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ (subset of $\mathbb{R}^d$ or a manifold)
- Moment characterisation:
  - Expectation $\mu(\mathbf{s}) = \mathbb{E}[u(\mathbf{s})]$
  - Covariance $\mathcal{R}(\mathbf{s}, \mathbf{s}') = \mathbb{C}[u(\mathbf{s}), u(\mathbf{s}')]$, symmetric positive definite function.
- Precision operator; inverse covariance: $\mathcal{Q} = \mathcal{R}^{-1}$

  In practice, easier conditions for valid models
- Reproducing Kernel Hilbert Space (RKHS) $H_\mathcal{Q}$: Inner product

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle f, \mathcal{Q} g \rangle_\mathcal{D}$$

  and squared norm $\|f\|^2 = \langle f, f \rangle_{H_\mathcal{Q}}$
- $m(\cdot) = \mathbb{E}(u(\cdot) - \mu(\cdot) | \{u(\mathbf{s}_k)\}) \in H_\mathcal{Q}$ but $u(\cdot) - \mu(\cdot) \notin H_\mathcal{Q}$; the process is less smooth!
- Spatial and spatio-temporal stochastic PDEs generate random field models:

$$\mathcal{L} u(\mathbf{s}) \, \mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s})$$
$$\mathcal{Q}_u = \mathcal{L}^* \mathcal{L}$$
$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle \mathcal{L} f, \mathcal{L} g \rangle_\mathcal{D}$$

  Can work directly with the precision or inner product; no need to know the covariance!

The eternal quest for spatial dependence models

- Gaussian random field: $u(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ (subset of $\mathbb{R}^d$ or a manifold)
- Moment characterisation:
  - Expectation $\mu(\mathbf{s}) = \mathbb{E}[u(\mathbf{s})]$
  - Covariance $\mathcal{R}(\mathbf{s}, \mathbf{s}') = \mathbb{C}[u(\mathbf{s}), u(\mathbf{s}')]$, symmetric positive definite function.
- Precision operator; inverse covariance: $\mathcal{Q} = \mathcal{R}^{-1}$

  In practice, easier conditions for valid models
- Reproducing Kernel Hilbert Space (RKHS) $H_\mathcal{Q}$: Inner product

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle f, \mathcal{Q} g \rangle_\mathcal{D}$$

  and squared norm $\|f\|^2 = \langle f, f \rangle_{H_\mathcal{Q}}$
- $m(\cdot) = \mathbb{E}(u(\cdot) - \mu(\cdot) | \{u(\mathbf{s}_k)\}) \in H_\mathcal{Q}$ but $u(\cdot) - \mu(\cdot) \notin H_\mathcal{Q}$; the process is less smooth!
- Spatial and spatio-temporal stochastic PDEs generate random field models:

$$\mathcal{L} u(\mathbf{s})\,\mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s})$$

$$\mathcal{Q}_u = \mathcal{L}^* \mathcal{L}$$

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle \mathcal{L} f, \mathcal{L} g \rangle_\mathcal{D}$$

  Can work directly with the precision or inner product; no need to know the covariance!

The eternal quest for spatial dependence models

- Gaussian random field: $u(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ (subset of $\mathbb{R}^d$ or a manifold)
- Moment characterisation:
  - Expectation $\mu(\mathbf{s}) = \mathbb{E}[u(\mathbf{s})]$
  - Covariance $\mathcal{R}(\mathbf{s}, \mathbf{s}') = \mathbb{C}[u(\mathbf{s}), u(\mathbf{s}')]$, symmetric positive definite function.
- Precision operator; inverse covariance: $\mathcal{Q} = \mathcal{R}^{-1}$

  In practice, easier conditions for valid models
- Reproducing Kernel Hilbert Space (RKHS) $H_\mathcal{Q}$: Inner product

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle f, \mathcal{Q} g \rangle_\mathcal{D}$$

  and squared norm $\|f\|^2 = \langle f, f \rangle_{H_\mathcal{Q}}$
- $m(\cdot) = \mathbb{E}(u(\cdot) - \mu(\cdot) | \{u(\mathbf{s}_k)\}) \in H_\mathcal{Q}$ but $u(\cdot) - \mu(\cdot) \notin H_\mathcal{Q}$; the process is less smooth!
- Spatial and spatio-temporal stochastic PDEs generate random field models:

$$\mathcal{L} u(\mathbf{s}) \,\mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s})$$

$$\mathcal{Q}_u = \mathcal{L}^* \mathcal{L}$$

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle \mathcal{L} f, \mathcal{L} g \rangle_\mathcal{D}$$

  Can work directly with the precision or inner product; no need to know the covariance!

The eternal quest for spatial dependence models

- Gaussian random field: $u(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ (subset of $\mathbb{R}^d$ or a manifold)
- Moment characterisation:
  - Expectation $\mu(\mathbf{s}) = \mathbb{E}[u(\mathbf{s})]$
  - Covariance $\mathcal{R}(\mathbf{s}, \mathbf{s}') = \mathbb{C}[u(\mathbf{s}), u(\mathbf{s}')]$, symmetric positive definite function.
- Precision operator; inverse covariance: $\mathcal{Q} = \mathcal{R}^{-1}$

  In practice, easier conditions for valid models
- Reproducing Kernel Hilbert Space (RKHS) $H_{\mathcal{Q}}$: Inner product

$$\langle f, g \rangle_{H_{\mathcal{Q}}} = \langle f, \mathcal{Q} g \rangle_{\mathcal{D}}$$

  and squared norm $\|f\|^2 = \langle f, f \rangle_{H_{\mathcal{Q}}}$
- $m(\cdot) = \mathbb{E}(u(\cdot) - \mu(\cdot) | \{u(\mathbf{s}_k)\}) \in H_{\mathcal{Q}}$ but $u(\cdot) - \mu(\cdot) \notin H_{\mathcal{Q}}$; the process is less smooth!
- Spatial and spatio-temporal stochastic PDEs generate random field models:

$$\mathcal{L} u(\mathbf{s}) \, \mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s})$$

$$\mathcal{Q}_u = \mathcal{L}^* \mathcal{L}$$

$$\langle f, g \rangle_{H_{\mathcal{Q}}} = \langle \mathcal{L} f, \mathcal{L} g \rangle_{\mathcal{D}}$$

  Can work directly with the precision or inner product; no need to know the covariance!

The eternal quest for spatial dependence models

- Gaussian random field: $u(\mathbf{s})$, $\mathbf{s} \in \mathcal{D}$ (subset of $\mathbb{R}^d$ or a manifold)
- Moment characterisation:
  - Expectation $\mu(\mathbf{s}) = \mathbb{E}[u(\mathbf{s})]$
  - Covariance $\mathcal{R}(\mathbf{s}, \mathbf{s}') = \mathbb{C}[u(\mathbf{s}), u(\mathbf{s}')]$, symmetric positive definite function.
- Precision operator; inverse covariance: $\mathcal{Q} = \mathcal{R}^{-1}$

  In practice, easier conditions for valid models
- Reproducing Kernel Hilbert Space (RKHS) $H_\mathcal{Q}$: Inner product

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle f, \mathcal{Q} g \rangle_\mathcal{D}$$

  and squared norm $\|f\|^2 = \langle f, f \rangle_{H_\mathcal{Q}}$
- $m(\cdot) = \mathbb{E}(u(\cdot) - \mu(\cdot) | \{u(\mathbf{s}_k)\}) \in H_\mathcal{Q}$ but $u(\cdot) - \mu(\cdot) \notin H_\mathcal{Q}$; the process is less smooth!
- Spatial and spatio-temporal stochastic PDEs generate random field models:

$$\mathcal{L} u(\mathbf{s}) \,\mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s})$$

$$\mathcal{Q}_u = \mathcal{L}^* \mathcal{L}$$

$$\langle f, g \rangle_{H_\mathcal{Q}} = \langle \mathcal{L} f, \mathcal{L} g \rangle_\mathcal{D}$$

  Can work directly with the precision or inner product; no need to know the covariance!

Direct Bayesian inference:
The inner core of the Integrated Nested Laplace Approximation method

- Latent Gaussian model structure

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}) \quad \text{(precision parameters)} \qquad \eta(\mathbf{s}) = \sum_{k=1}^{n} \psi(\mathbf{s}) u_k \quad \text{(predictor)}$$

$$\boldsymbol{u}|\boldsymbol{\theta} \sim \mathsf{N}[\boldsymbol{\mu}_u, \boldsymbol{Q}_u^{-1}] \quad \text{(latent field)} \qquad \boldsymbol{y}|\boldsymbol{\theta},\boldsymbol{u} \sim p(\boldsymbol{y}|\boldsymbol{\theta},\eta) \quad \text{(observations)}$$

- Conditional log-posterior mode ($\boldsymbol{\mu}_{u|y}$) and Hessian ($\boldsymbol{Q}_{u|y}$), for each $\boldsymbol{\theta}$, by iteration:

$$\boldsymbol{g}_y^* = -\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{u}} \log p(\boldsymbol{y}|\boldsymbol{\theta},\eta)\Big|_{\boldsymbol{u}=\boldsymbol{u}^*}$$

$$\boldsymbol{H}_y^* = -\frac{\mathrm{d}^2}{\mathrm{d}\boldsymbol{u}\mathrm{d}\boldsymbol{u}^\top} \log p(\boldsymbol{y}|\boldsymbol{\theta},\eta)\Big|_{\boldsymbol{u}=\boldsymbol{u}^*}$$

$$\boldsymbol{Q}_{u|y} = \boldsymbol{Q}_u + \boldsymbol{H}_y^*$$

$$\boldsymbol{Q}_{u|y}(\boldsymbol{\mu}_{u|y} - \boldsymbol{\mu}_u) = \boldsymbol{H}_u^*(\boldsymbol{u}^* - \boldsymbol{\mu}_u) - \boldsymbol{g}_y^*$$

Direct Bayesian inference: The inner core of the Integrated Nested Laplace Approximation method

- Latent Gaussian model structure

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}) \quad \text{(precision parameters)} \qquad \eta(\mathbf{s}) = \sum_{k=1}^{n} \psi(\mathbf{s}) u_k \quad \text{(predictor)}$$

$$\boldsymbol{u}|\boldsymbol{\theta} \sim \mathsf{N}[\boldsymbol{\mu}_u, \boldsymbol{Q}_u^{-1}] \quad \text{(latent field)} \qquad \boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{u} \sim p(\boldsymbol{y}|\boldsymbol{\theta}, \eta) \quad \text{(observations)}$$

- Conditional log-posterior mode ($\boldsymbol{\mu}_{u|y}$) and Hessian ($\boldsymbol{Q}_{u|y}$), for each $\boldsymbol{\theta}$, by iteration:

$$\boldsymbol{g}_y^* = - \frac{\mathrm{d}}{\mathrm{d}\boldsymbol{u}} \log p(\boldsymbol{y}|\boldsymbol{\theta}, \eta) \Big|_{\boldsymbol{u}=\boldsymbol{u}^*}$$

$$\boldsymbol{H}_y^* = - \frac{\mathrm{d}^2}{\mathrm{d}\boldsymbol{u}\mathrm{d}\boldsymbol{u}^\top} \log p(\boldsymbol{y}|\boldsymbol{\theta}, \eta) \Big|_{\boldsymbol{u}=\boldsymbol{u}^*}$$

$$\boldsymbol{Q}_{u|y} = \boldsymbol{Q}_u + \boldsymbol{H}_y^*$$

$$\boldsymbol{Q}_{u|y}(\boldsymbol{\mu}_{u|y} - \boldsymbol{\mu}_u) = \boldsymbol{H}_u^*(\boldsymbol{u}^* - \boldsymbol{\mu}_u) - \boldsymbol{g}_y^*$$

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?
Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?
Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$ but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$ and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?
Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?

Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?
Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure?

Question 2: Is the "separable" vs "non-separable" dichotomy sufficient?

Spatio-temporal separability for functions, covariances, and precisions

- Functional separability for $\mathbf{s} \in \mathcal{D}$ and $t \in \mathcal{T}$
  - Addition: $w(\mathbf{s}, t) = u(\mathbf{s}) + v(t)$
  - Multiplication $w(\mathbf{s}, t) = u(\mathbf{s})v(t)$ (degrees of freedom $|\mathcal{D}| + |\mathcal{T}|$)
- Covariance separability
  - Addition: $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}') + \mathcal{R}_v(t, t')$
  - Multiplication $\mathcal{R}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{R}_u(\mathbf{s}, \mathbf{s}')\mathcal{R}_v(t, t')$ (degrees of freedom $|\mathcal{D}| \cdot |\mathcal{T}|$)
  - Simple to construct, but with some unrealistic properties
  - Additive combination: $\sum_k \mathcal{R}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{R}_{v_k}(t, t')$ (sum of separable processes)
- "Precision separability"
  - Addition: $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}') + \mathcal{Q}_v(t, t')$
    but $\mathcal{R}_w = \mathcal{Q}_w^{-1} = (\mathcal{Q}_u + \mathcal{Q}_v)^{-1}$ (not covariance separable!)
  - Multiplication $\mathcal{Q}_w[(\mathbf{s}, t), (\mathbf{s}', t')] = \mathcal{Q}_u(\mathbf{s}, \mathbf{s}')\mathcal{Q}_v(t, t')$
    and $\mathcal{R}_w = \mathcal{Q}_w^{-1} = \mathcal{Q}_u^{-1}\mathcal{Q}_v^{-1}$ (covariance separable!)
  - Additive combination: $\sum_k \mathcal{Q}_{u_k}(\mathbf{s}, \mathbf{s}')\mathcal{Q}_{v_k}(t, t')$ (what is this!?!)

Question 1: Are there interpretable process constructions that lead to this structure? Yes!
Question 2: Is the "separable" vs "non-separable" dichotomy sufficient? No!
Non-stationarity and asymmetry also important.

From temporal random walks to spatio-temporal diffusion

- Spatial Whittle-Matérn models with $\mathcal{L}_s = \gamma_s^2 - \Delta$:

$$\mathcal{L}_s^{\alpha_s/2} u(\mathbf{s})\,\mathrm{d}\mathbf{s} = \mathrm{d}\mathcal{W}(\mathbf{s}) \quad \text{(spatial white noise)}$$

  Precision $\mathcal{Q} = \mathcal{L}_s^{\alpha_s}$, Matérn covariance on $\mathbb{R}^d$.

- Separable space-time model (separable vector Ornstein-Uhlenbeck/AR(1) process):

$$\left(\frac{\partial}{\partial t} + \kappa\right) \mathcal{L}_s^{\alpha_s/2} u(\mathbf{s}, t)\,\mathrm{d}\mathbf{s}\,\mathrm{d}t = \mathrm{d}\mathcal{W}(\mathbf{s}, t) \quad \text{(spatio-temporal white noise)}$$

  Precision $\mathcal{Q} = \left(\kappa^2 - \frac{\partial^2}{\partial t^2}\right) \mathcal{L}_s^{\alpha_s}$, covariance is a product of a temporal Matérn kernel and the spatial covariance.

Prediction

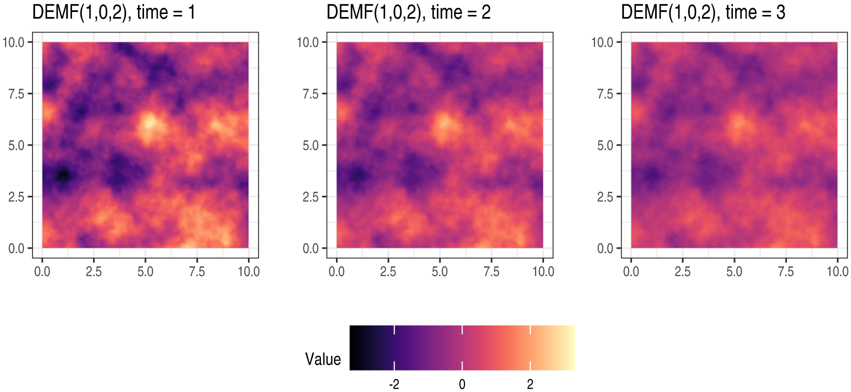

Conditional expectations into the future decay pointwise towards zero; no spatial dynamics.

Diffusion extension of Matérn fields (DEMF)

- Non-separable space-time $\text{DEMF}(\alpha_t, \alpha_s, \alpha_e)$ model for $(\mathbf{s}, t) \in \mathcal{D} \times \mathcal{T}$:

$$\gamma_e \mathcal{L}_s^{\alpha_e/2} \left(\gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s/2} \right)^{\alpha_t} u(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t \overset{d}{=} \gamma_e \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t^2 \frac{\partial^2}{\partial t^2} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2} u(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t = \mathrm{d}\mathcal{W}(\mathbf{s}, t),$$

where $\gamma_e, \gamma_t > 0$, and $\alpha_t > 0$, $\alpha_s, \alpha_e \geq 0$.

- In the stationary case, the resulting field has Matérn covariance for every time point
- The spatial smoothness is $\nu_s = \alpha_s(\alpha_t - 1/2) + \alpha_e - d/2$
- The temporal smoothness is $\nu_t = \min[\alpha_t - 1/2, \nu_s/\alpha_s]$.
- Non-separability parameter: $\beta_s = 1 - \frac{\alpha_e}{\nu_s + d/2} \in [0, 1]$
- Tensor product basis discretisation for integer α_t gives precision matrix structure

$$\boldsymbol{Q} = \gamma_e^2 \sum_{k=0}^{2\alpha_t} \gamma_t^k \boldsymbol{J}_{\alpha_t, k/2} \otimes \boldsymbol{K}_{\alpha_s(\alpha_t - k/2) + \alpha_e}$$

where $\boldsymbol{J}_{\cdot,\cdot}$ are purely temporal and $\boldsymbol{K}_{\cdot}$ are purely spatial.
This is what we were looking for!

Diffusion extension of Matérn fields (DEMF)

- Non-separable space-time DEMF$(\alpha_t, \alpha_s, \alpha_e)$ model for $(\mathbf{s}, t) \in \mathcal{D} \times \mathcal{T}$:

$$\gamma_e \mathcal{L}_s^{\alpha_e/2} \left(\gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s/2} \right)^{\alpha_t} u(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t \overset{d}{=} \gamma_e \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t^2 \frac{\partial^2}{\partial t^2} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2} u(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t = \mathrm{d}\mathcal{W}(\mathbf{s}, t),$$

  where $\gamma_e, \gamma_t > 0$, and $\alpha_t > 0$, $\alpha_s, \alpha_e \geq 0$.
- In the stationary case, the resulting field has Matérn covariance for every time point
- The spatial smoothness is $\nu_s = \alpha_s(\alpha_t - 1/2) + \alpha_e - d/2$
- The temporal smoothness is $\nu_t = \min[\alpha_t - 1/2, \nu_s/\alpha_s]$.
- Non-separability parameter: $\beta_s = 1 - \frac{\alpha_e}{\nu_s + d/2} \in [0, 1]$
- Tensor product basis discretisation for integer α_t gives precision matrix structure

$$\boldsymbol{Q} = \gamma_e^2 \sum_{k=0}^{2\alpha_t} \gamma_t^k \boldsymbol{J}_{\alpha_t, k/2} \otimes \boldsymbol{K}_{\alpha_s(\alpha_t - k/2) + \alpha_e}$$

  where $\boldsymbol{J}_{\cdot,\cdot}$ are purely temporal and $\boldsymbol{K}_{\cdot}$ are purely spatial.
  This is what we were looking for!

Diffusion extension of Matérn fields (DEMF)

- Non-separable space-time DEMF$(\alpha_t, \alpha_s, \alpha_e)$ model for $(\mathbf{s}, t) \in \mathcal{D} \times \mathcal{T}$:

$$\gamma_e \mathcal{L}_s^{\alpha_e/2} \left(\gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s/2} \right)^{\alpha_t} u(\mathbf{s}, t) \,\mathrm{d}\mathbf{s}\,\mathrm{d}t \overset{d}{=} \gamma_e \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t^2 \frac{\partial^2}{\partial t^2} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2} u(\mathbf{s}, t) \,\mathrm{d}\mathbf{s}\,\mathrm{d}t = \mathrm{d}\mathcal{W}(\mathbf{s}, t),$$

  where $\gamma_e, \gamma_t > 0$, and $\alpha_t > 0$, $\alpha_s, \alpha_e \geq 0$.
- In the stationary case, the resulting field has Matérn covariance for every time point
- The spatial smoothness is $\nu_s = \alpha_s(\alpha_t - 1/2) + \alpha_e - d/2$
- The temporal smoothness is $\nu_t = \min[\alpha_t - 1/2, \nu_s/\alpha_s]$.
- Non-separability parameter: $\beta_s = 1 - \frac{\alpha_e}{\nu_s + d/2} \in [0, 1]$
- Tensor product basis discretisation for integer α_t gives precision matrix structure

$$\boldsymbol{Q} = \gamma_e^2 \sum_{k=0}^{2\alpha_t} \gamma_t^k \boldsymbol{J}_{\alpha_t, k/2} \otimes \boldsymbol{K}_{\alpha_s(\alpha_t - k/2) + \alpha_e}$$

  where $\boldsymbol{J}_{\cdot,\cdot}$ are purely temporal and $\boldsymbol{K}_{\cdot}$ are purely spatial.
  This is what we were looking for!

Non-separable covariances, from spectral inversion; $\mathbb{R}^2 \times \mathbb{R}$

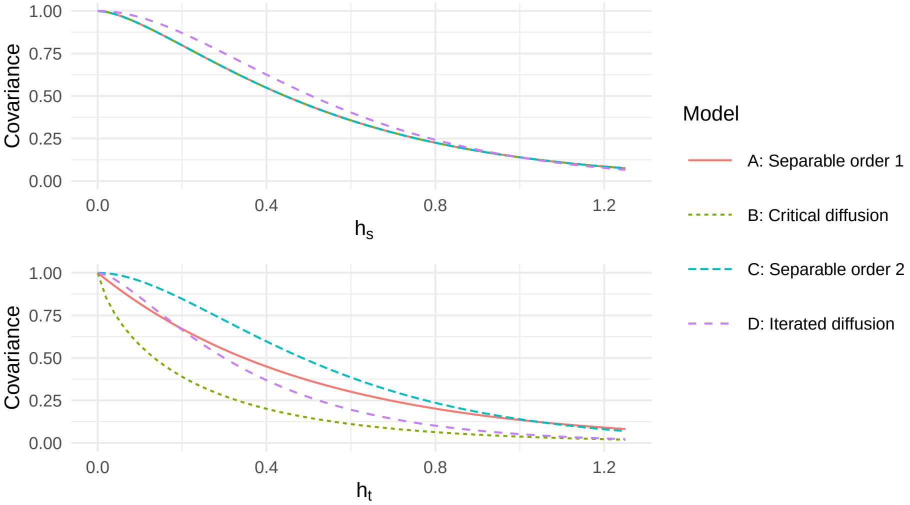

Non-separable covariances, from spectral inversion; $\mathbb{R}^2 \times \mathbb{R}$

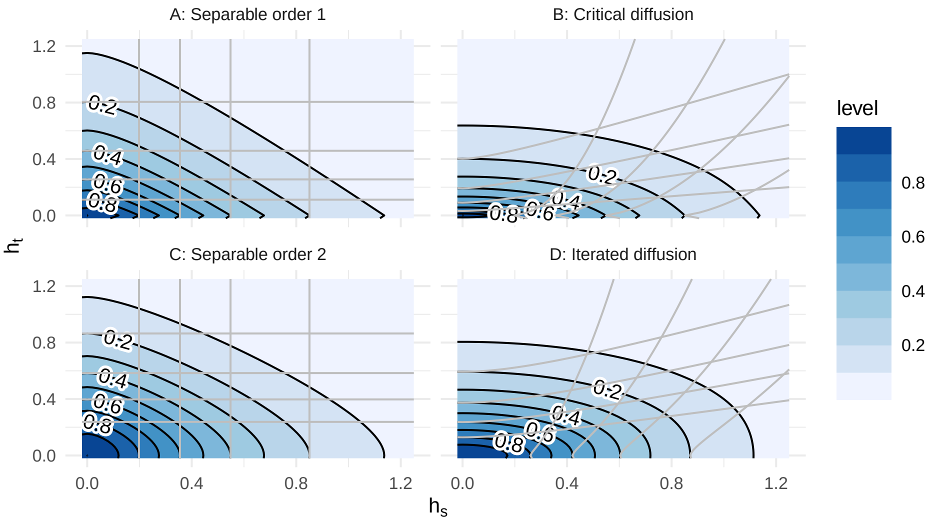

Prediction

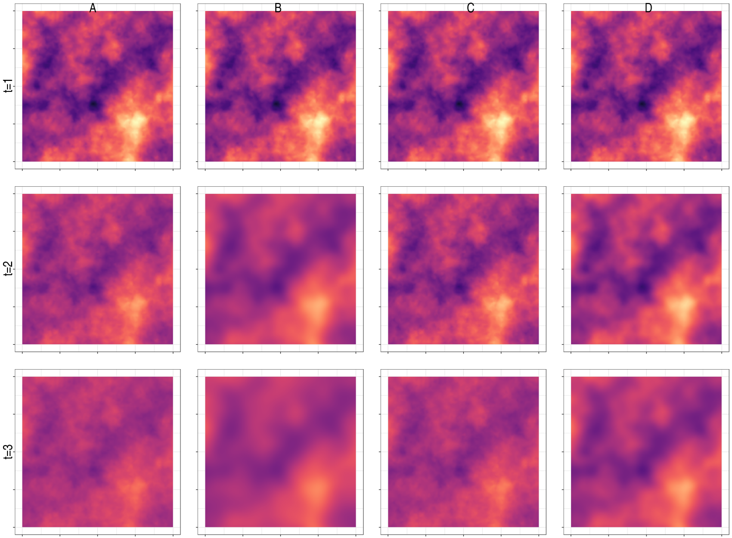

Manifold processes and metric change

- When $\mathcal{D}$ is a curved manifold, e.g. $\mathbb{S}^2$, the SPDE formulation of Matérn fields defines generalised Whittle-Matérn fields
- Stationarity only definable and achievable for manifolds with constant curvature
- For compact manifolds: implicit spectral representation based on eigenfunctions of the Laplacian; spherical harmonics on $\mathbb{S}^2$.

Metric change:

- Change of variables between two manifold domains $\widetilde{\mathcal{D}}$ and $\mathcal{D}$ retains marginal variances
- Let $r(\mathbf{s})$ be the geometric mean metric stretching and anisotropy $\overline{\boldsymbol{H}}(\mathbf{s})$ with $\det \overline{\boldsymbol{H}}(\mathbf{s}) \equiv 1$.Then

$$(1 - \widetilde{\nabla} \cdot \widetilde{\nabla})\widetilde{u}(\mathbf{s})\,\mathrm{d}\mathbf{s} = \widetilde{\mathrm{d}\mathcal{W}}(\mathbf{s}), \quad \mathbf{s} \in \widetilde{\mathcal{D}}$$

$$\frac{1}{r(\mathbf{s})^d}(1 - r(\mathbf{s})^d \nabla \cdot r(\mathbf{s})^{2-d}\overline{\boldsymbol{H}}(\mathbf{s})\nabla)u(\mathbf{s})\,\mathrm{d}\mathbf{s} = \frac{1}{r(\mathbf{s})^{d/2}}\mathrm{d}\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}$$

- For *given* $r(\mathbf{s})$ and $\overline{\boldsymbol{H}}(\mathbf{s})$, the $u(\mathbf{s})$ process can have near constant variance, despite $\widetilde{\mathcal{D}}$ having non-constant curvature.

Manifold processes and metric change

- When $\mathcal{D}$ is a curved manifold, e.g. $\mathbb{S}^2$, the SPDE formulation of Matérn fields defines generalised Whittle-Matérn fields
- Stationarity only definable and achievable for manifolds with constant curvature
- For compact manifolds: implicit spectral representation based on eigenfunctions of the Laplacian; spherical harmonics on $\mathbb{S}^2$.

Metric change:

- Change of variables between two manifold domains $\widetilde{\mathcal{D}}$ and $\mathcal{D}$ retains marginal variances
- Let $r(\mathbf{s})$ be the geometric mean metric stretching and anisotropy $\overline{\boldsymbol{H}}(\mathbf{s})$ with $\det \overline{\boldsymbol{H}}(\mathbf{s}) \equiv 1$. Then

$$(1 - \widetilde{\nabla} \cdot \widetilde{\nabla}) \widetilde{u}(\mathbf{s}) \, \mathrm{d}\mathbf{s} = \widetilde{\mathrm{d}\mathcal{W}}(\mathbf{s}), \quad \mathbf{s} \in \widetilde{\mathcal{D}}$$

$$\frac{1}{r(\mathbf{s})^d} (1 - r(\mathbf{s})^d \nabla \cdot r(\mathbf{s})^{2-d} \overline{\boldsymbol{H}}(\mathbf{s}) \nabla) u(\mathbf{s}) \, \mathrm{d}\mathbf{s} = \frac{1}{r(\mathbf{s})^{d/2}} \mathrm{d}\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}$$

- For *given* $r(\mathbf{s})$ and $\overline{\boldsymbol{H}}(\mathbf{s})$, the $u(\mathbf{s})$ process can have near constant variance, despite $\widetilde{\mathcal{D}}$ having non-constant curvature.

Manifold processes and metric change

- When $\mathcal{D}$ is a curved manifold, e.g. $\mathbb{S}^2$, the SPDE formulation of Matérn fields defines generalised Whittle-Matérn fields
- Stationarity only definable and achievable for manifolds with constant curvature
- For compact manifolds: implicit spectral representation based on eigenfunctions of the Laplacian; spherical harmonics on $\mathbb{S}^2$.

Metric change:

- Change of variables between two manifold domains $\widetilde{\mathcal{D}}$ and $\mathcal{D}$ retains marginal variances
- Let $r(\mathbf{s})$ be the geometric mean metric stretching and anisotropy $\overline{\boldsymbol{H}}(\mathbf{s})$ with $\det \overline{\boldsymbol{H}}(\mathbf{s}) \equiv 1$.Then

$$
(1 - \widetilde{\nabla} \cdot \widetilde{\nabla})\widetilde{u}(\mathbf{s})\,\mathrm{d}\mathbf{s} = \widetilde{\mathrm{d}\mathcal{W}}(\mathbf{s}), \quad \mathbf{s} \in \widetilde{\mathcal{D}}
$$

$$
\frac{1}{r(\mathbf{s})^d}(1 - r(\mathbf{s})^d \nabla \cdot r(\mathbf{s})^{2-d}\overline{\boldsymbol{H}}(\mathbf{s})\nabla)u(\mathbf{s})\,\mathrm{d}\mathbf{s} = \frac{1}{r(\mathbf{s})^{d/2}}\mathrm{d}\mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}
$$

- For *given* $r(\mathbf{s})$ and $\overline{\boldsymbol{H}}(\mathbf{s})$, the $u(\mathbf{s})$ process can have near constant variance, despite $\widetilde{\mathcal{D}}$ having non-constant curvature.

Separable non-stationarity in non-separable space-time

- The metric change technique can be used to define space-time separable non-stationarity within a non-separable space-time model
- Let $\mathcal{L}_{\mathbf{s}} = 1 - r_{\mathbf{s}}^d \nabla \cdot r_{\mathbf{s}}^{2-d} \overline{\boldsymbol{H}}_{\mathbf{s}} \nabla$, depending on location, $\mathbf{s} \in \mathcal{D}$
- Let γ_t be the desired local temporal range, depending on time, $t \in \mathbb{R}$
- Metric change:

$$(1 - \widetilde{\nabla} \cdot \widetilde{\nabla})^{\alpha_e/2} \left(-\frac{\partial^2}{\partial t^2} + (1 - \widetilde{\nabla} \cdot \widetilde{\nabla})^{\alpha_s} \right)^{\alpha_t/2} \widetilde{u}(\mathbf{s}, t) \,\mathrm{d}\mathbf{s}\,\mathrm{d}t = \mathrm{d}\widetilde{\mathcal{W}}(\mathbf{s}, t), \quad (\mathbf{s}, t) \in \widetilde{\mathcal{D}} \times \mathbb{R}$$

$$\frac{1}{\gamma_t r_{\mathbf{s}}^d} \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t \frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2} u(\mathbf{s}, t) \,\mathrm{d}\mathbf{s}\,\mathrm{d}t = \frac{1}{\gamma_t^{1/2} r_{\mathbf{s}}^{d/2}} \mathrm{d}\mathcal{W}(\mathbf{s}, t), \quad (\mathbf{s}, t) \in \mathcal{D} \times \mathbb{R}$$

- Precision has sum-product structure, but the parameter fields enter non-commutatively:

$$\mathcal{Q} = \left(-\frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} \gamma_t + (\mathcal{L}_s^*)^{\alpha_s} \right)^{\alpha_t/2} (\mathcal{L}_s^*)^{\alpha_e/2} \frac{1}{\gamma_t r_{\mathbf{s}}^d} \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t \frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2}$$

Separable non-stationarity in non-separable space-time

- The metric change technique can be used to define space-time separable non-stationarity within a non-separable space-time model
- Let $\mathcal{L}_{\mathbf{s}} = 1 - r_{\mathbf{s}}^d \nabla \cdot r_{\mathbf{s}}^{2-d} \overline{\boldsymbol{H}}_{\mathbf{s}} \nabla$, depending on location, $\mathbf{s} \in \mathcal{D}$
- Let γ_t be the desired local temporal range, depending on time, $t \in \mathbb{R}$
- Metric change:

$$(1 - \widetilde{\nabla} \cdot \widetilde{\nabla})^{\alpha_e/2} \left(-\frac{\partial^2}{\partial t^2} + (1 - \widetilde{\nabla} \cdot \widetilde{\nabla})^{\alpha_s} \right)^{\alpha_t/2} \widetilde{u}(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t = \widetilde{\mathrm{d}\mathcal{W}}(\mathbf{s}, t), \quad (\mathbf{s}, t) \in \widetilde{\mathcal{D}} \times \mathbb{R}$$

$$\frac{1}{\gamma_t r_{\mathbf{s}}^d} \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t \frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2} u(\mathbf{s}, t) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t = \frac{1}{\gamma_t^{1/2} r_{\mathbf{s}}^{d/2}} \mathrm{d}\mathcal{W}(\mathbf{s}, t), \quad (\mathbf{s}, t) \in \mathcal{D} \times \mathbb{R}$$

- Precision has sum-product structure, but the parameter fields enter non-commutatively:

$$\mathcal{Q} = \left(-\frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} \gamma_t + (\mathcal{L}_s^*)^{\alpha_s} \right)^{\alpha_t/2} (\mathcal{L}_s^*)^{\alpha_e/2} \frac{1}{\gamma_t r_{\mathbf{s}}^d} \mathcal{L}_s^{\alpha_e/2} \left(-\gamma_t \frac{\partial}{\partial t} \gamma_t \frac{\partial}{\partial t} + \mathcal{L}_s^{\alpha_s} \right)^{\alpha_t/2}$$

Basic temperature modelling on a daily timescale

- $\mu(s,t) = \beta_0 + \sum_k B_k(\mathbf{s})\beta_k(t) + \text{elevation}(\mathbf{s})\beta_e + u(\mathbf{s},t)$
- $y_i \sim \mathsf{N}(\mu(\mathbf{s}_i,t_i),\sigma_y^2)$
- Estimated with the four previous models for $u(\mathbf{s},t)$;
  separable and non-separable, temporal order 1 and 2.
- 1,649 spatial mesh nodes, 365 temporal nodes, 13,704 weather stations
- The non-separable models have slightly lower MSE&CRPS scores. However...
- Point estimates and leave-one-out predictions insufficient to clearly distinguish the models
- Non-separability matters more for spatio-temporal prediction and uncertainty quantification than for field reconstruction

Basic temperature modelling on a daily timescale

- $\mu(s,t) = \beta_0 + \sum_k B_k(\mathbf{s})\beta_k(t) + \text{elevation}(\mathbf{s})\beta_e + u(\mathbf{s},t)$
- $y_i \sim \mathsf{N}(\mu(\mathbf{s}_i, t_i), \sigma_y^2)$
- Estimated with the four previous models for $u(\mathbf{s},t)$;
  separable and non-separable, temporal order 1 and 2.
- 1,649 spatial mesh nodes, 365 temporal nodes, 13,704 weather stations
- The non-separable models have slightly lower MSE&CRPS scores. However...
- Point estimates and leave-one-out predictions insufficient to clearly distinguish the models
- Non-separability matters more for spatio-temporal prediction and uncertainty quantification than for field reconstruction

Seasonal pattern

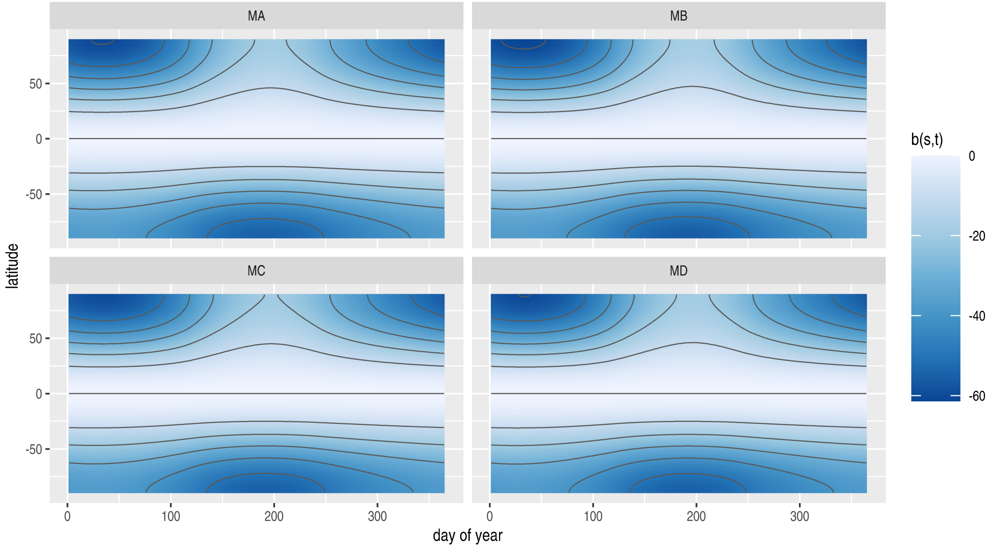

Non-separable component, January 1-3 2021

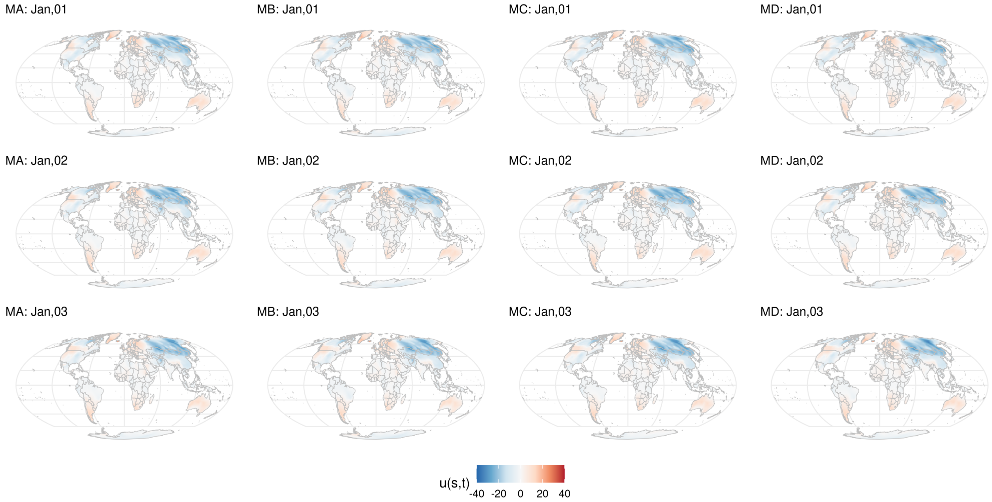

Non-separable component, January 4-6 2021

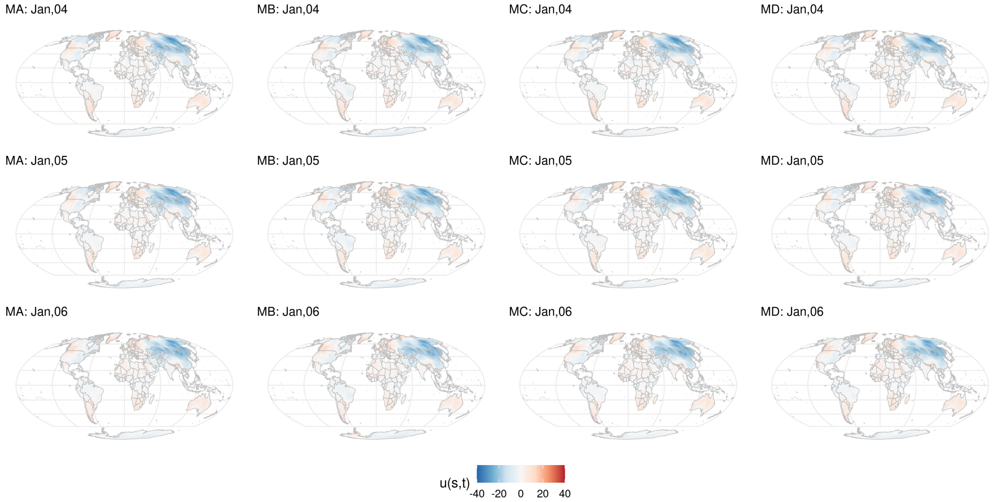

Non-separable component, January 7-9 2021

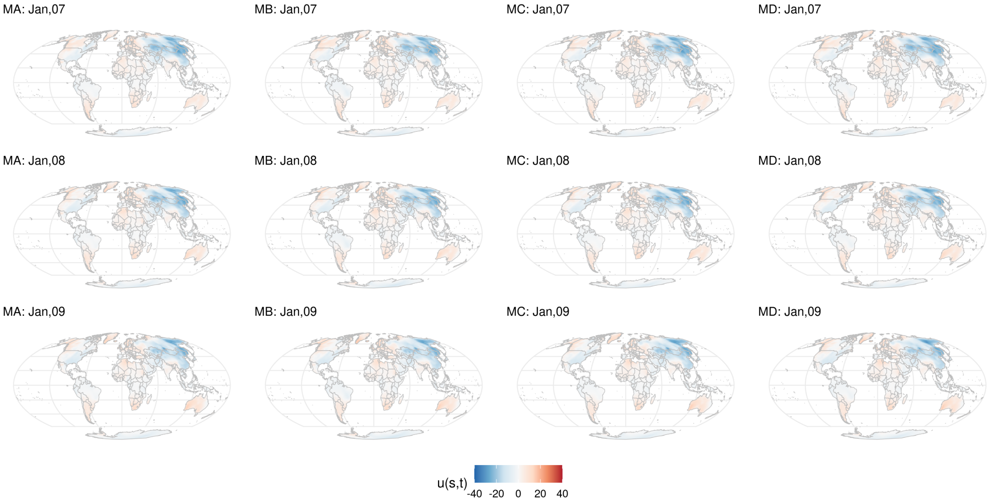

Non-separable component, January 10-12 2021

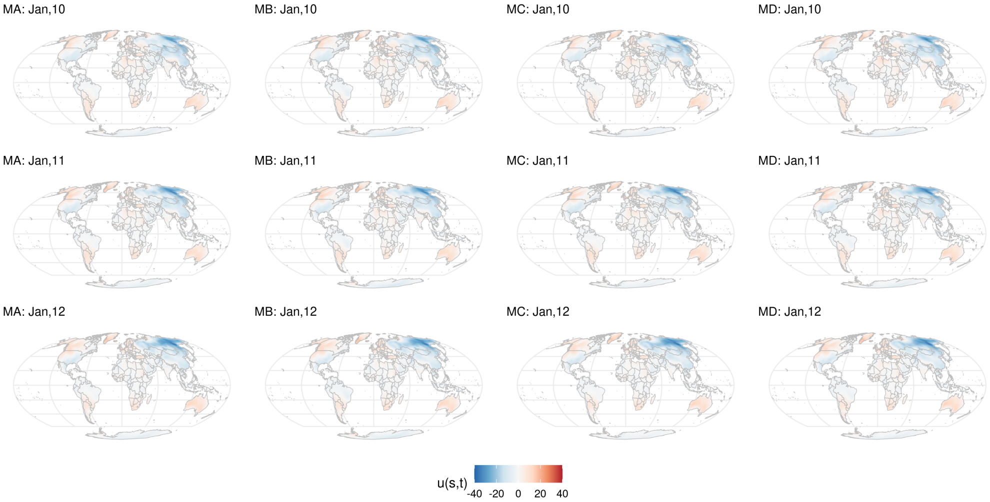

Summary

- Non-separability and non-stationarity are distinct concepts; both needed
- (Relatively) simple stochastic PDE concepts provide useful building blocks
- Mainfold domains and non-stationarity is theoreticaly simple
- Computational methods need to handle hierarchical structures, not just additive noise
- The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running (Lindgren et al, 2022, Spatial Statistics)
  https://arxiv.org/abs/2111.01084
  https://doi.org/10.1016/j.spasta.2022.100599
- A diffusion-based spatio-temporal extension of Gaussian Matérn fields (Lindgren et al)
  https://arxiv.org/abs/2006.04917
- inlabru documentation and examples:
  https://inlabru-org.github.io/inlabru/
- INLAspacetime implementation:
  https://github.com/eliaskrainski/INLAspacetime/

Summary

- Non-separability and non-stationarity are distinct concepts; both needed
- (Relatively) simple stochastic PDE concepts provide useful building blocks
- Mainfold domains and non-stationarity is theoreticaly simple
- Computational methods need to handle hierarchical structures, not just additive noise
- The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running (Lindgren et al, 2022, Spatial Statistics)
  https://arxiv.org/abs/2111.01084
  https://doi.org/10.1016/j.spasta.2022.100599
- A diffusion-based spatio-temporal extension of Gaussian Matérn fields (Lindgren et al)
  https://arxiv.org/abs/2006.04917
- `inlabru` documentation and examples:
  https://inlabru-org.github.io/inlabru/
- `INLAspacetime` implementation:
  https://github.com/eliaskrainski/INLAspacetime/