A diffusion-based spatio-temporal extension of Gaussian Matérn fields

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

Haakon Bakka, David Bolin, Elias Krainski, Haavard Rue

THE UNIVERSITY *of* EDINBURGH

Barcelona, 2024-04-15

Context: spatial and spatio-temporal modelling and estimation

Problem examples:

- Reconstructing past temperatures from weather stations, satellite, and ship measurements
- Estimating the abundance of dolphins, from ship observations
- Tracking electrical signals across the heart
- Estimating the habitats and movement of birds, from GPS measurements

Modelling and estimation tools:

- Additive models (GAMM/GLM/GLMM/LGM/etc)
- Stochastic processes, specifically Gaussian, dynamical and static
- Bayesian methods (MCMC, INLA, variational Bayes)
- Software for special cases, and for general models

Context: spatial and spatio-temporal modelling and estimation

Problem examples:

- Reconstructing past temperatures from weather stations, satellite, and ship measurements
- Estimating the abundance of dolphins, from ship observations
- Tracking electrical signals across the heart
- Estimating the habitats and movement of birds, from GPS measurements

Modelling and estimation tools:

- Additive models (GAMM/GLM/GLMM/LGM/etc)
- Stochastic processes, specifically Gaussian, dynamical and static
- Bayesian methods (MCMC, INLA, variational Bayes)
- Software for special cases, and for general models

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

- Latent Gaussian model structure (Bayesian GAMs with Gaussian process components)

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

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

- 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 (Bayesian GAMs with Gaussian process components)

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

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

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

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

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

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

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

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 such as $\mathbb{S}^2$)
- 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 such as $\mathbb{S}^2$)
- 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 such as $\mathbb{S}^2$)
- 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 such as $\mathbb{S}^2$)
- 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 such as $\mathbb{S}^2$)
- 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!

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 cov-product-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 cov-product-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 cov-product-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 cov-product-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 cov-product-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 cov-product-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 cov-product-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!

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 for $u(\mathbf{s})$ on $\mathbb{R}^d$.

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

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

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

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 for $u(\mathbf{s})$ on $\mathbb{R}^d$.

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

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

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

Prediction

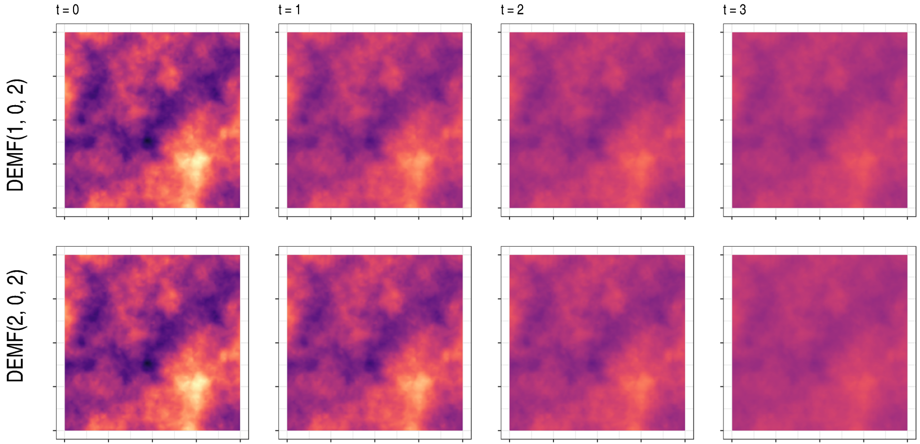

Conditional expectations into the future decay pointwise towards zero; no spatial dynamics.

Diffusion extension of Matérn fields (DEMF)

- Non-separable space-time 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
- 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
- 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
- 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!

Spectra and finite element structure

- Fourier spectra are based on eigenfunctions $e_{\boldsymbol{\omega}}(\mathbf{s})$ of $-\Delta$.
  On $\mathbb{R}^d$, $-\Delta e_{\boldsymbol{\lambda}}(\mathbf{s}) = \|\boldsymbol{\lambda}\|^2 e_{\boldsymbol{\lambda}}(\mathbf{s})$, and $e_{\boldsymbol{\lambda}}(\mathbf{s})$ are harmonic functions.
- The stationary spectrum on $\mathbb{R}^d \times \mathbb{R}$ is

$$\widehat{\mathcal{R}}(\boldsymbol{\lambda}, \omega) = \frac{1}{(2\pi)^{d+1}\tau^2(\kappa^2 + \|\boldsymbol{\lambda}\|^2)^{\alpha_e}\left[\phi^2\omega^2 + (\kappa^2 + \|\boldsymbol{\lambda}\|^2)^{\alpha_s}\right]^{\alpha_t}}$$

- On $\mathbb{S}^2$, $-\Delta e_k(\mathbf{s}) = \lambda_k e_k(\mathbf{s}) = k(k+1)e_k(\mathbf{s})$, and e_k are spherical harmonics.
- The isotropic spectrum on $\mathbb{S}^2 \times \mathbb{R}$ is

$$\widehat{\mathcal{R}}(k, \omega) \propto \frac{2k+1}{\tau^2(\kappa^2 + \lambda_k)^{\alpha_e}\left[\phi^2\omega^2 + (\kappa^2 + \lambda_k)^{\alpha_s}\right]^{\alpha_t}}$$

- The finite element approximation has structure

$$u(\boldsymbol{s}, t) = \sum_{i,j} \psi_i^{[s]}(\boldsymbol{s})\psi_j^{[t]}(t) x_{ij}, \quad \boldsymbol{x} \sim \mathsf{N}(\mathbf{0}, \boldsymbol{Q}^{-1}), \quad \boldsymbol{Q} = \sum_{k=0}^{\alpha_t+\alpha_s+\alpha_e} \boldsymbol{M}_k^{[t]} \otimes \boldsymbol{M}_k^{[\mathbf{s}]}$$

  even, e.g., if the spatial scale parameter κ is spatially varying.

Non-separable space-time: Matérn driven heat equation

The iterated heat equation is a simple non-separable space-time SPDE family:

$$\left[\phi\frac{\partial}{\partial t} + (\kappa^2 - \Delta)^{\alpha_s/2}\right]^{\alpha_t} u(\mathbf{s}, t)\,\mathrm{d}t = \mathrm{d}\mathcal{E}_{(\kappa^2-\Delta)^{\alpha_e}}(\mathbf{s}, t)/\tau$$

For constant parameters, $u(\mathbf{s}, t)$ has spatial Matérn covariance (for each t) on $\mathbb{R}^2$ and a generalised Matérn-Whittle sense on $\mathbb{S}^2$.
Smoothness properties (can be derived from the spectra):

$$\begin{cases} \nu_t = \min\left[\alpha_t - \frac{1}{2}, \frac{\nu_s}{\alpha_s}\right], \\ \nu_s = \alpha_e + \alpha_s\left(\alpha_t - \frac{1}{2}\right) - \frac{d}{2}, \\ \beta_s = 1 - \frac{\alpha_e}{\nu_s + d/2}, \end{cases} \qquad \begin{cases} \alpha_t = \nu_t \max\left(1, \frac{\beta_s}{\beta_*(\nu_s, d)}\right) + \frac{1}{2}, \\ \alpha_s = \frac{\nu_s}{\nu_t}\min\left(\frac{\beta_s}{\beta_*(\nu_s, d)}, 1\right) = \frac{1}{\nu_t}\min\left[(\nu_s + d/2)\beta_s, \nu_s\right], \\ \alpha_e = \frac{1 - \beta_s}{\beta_*(\nu_s, d)}\nu_s = (\nu_s + d/2)(1 - \beta_s), \end{cases}$$

where $\beta_*(\nu_s, d) = \frac{\nu_s}{\nu_s + d/2}$, and $\beta_s \in [0, 1]$ is a non-separability parameter.

Smoothness properties

α_t	α_s	α_e	Type	ν_t	ν_s
α_t	α_s	α_e	General	$\min\left[\alpha_t - \frac{1}{2}, \frac{\nu_s}{\alpha_s}\right]$	$\alpha_e + \alpha_s(\alpha_t - \frac{1}{2}) - \frac{d}{2}$
α_t	0	α_e	Separable	$\alpha_t - \frac{1}{2}$	$\alpha_e - \frac{d}{2}$
α_t	α_s	$\frac{d}{2}$	Critical	$\alpha_t - \frac{1}{2}$	$\alpha_s(\alpha_t - \frac{1}{2})$
α_t	α_s	0	Fully non-separable	$\alpha_t - \frac{1}{2} - \frac{d}{2\alpha_s}$	$\alpha_s(\alpha_t - \frac{1}{2}) - \frac{d}{2}$
1	2	$\alpha_e > \frac{d}{2}$	Sub-critical diffusion	$1/2$	$\alpha_e + 1 - \frac{d}{2}$
1	2	$\frac{d}{2}$	Critical diffusion	$1/2$	1
1	2	$\frac{d}{2} - 1 < \alpha_e < \frac{d}{2}$	Super-critical diffusion	$\nu_s/2$	$\alpha_e + 1 - \frac{d}{2}$
1	0	2	Separable	$1/2$	$2 - \frac{d}{2}$
$3/2$	2	0	Fractional diffusion	$1 - \frac{d}{4}$	$2 - \frac{d}{2}$
2	2	0	Iterated diffusion	$\frac{3}{2} - \frac{d}{4}$	$3 - \frac{d}{2}$

Four specific DEMF models on $\mathcal{R}^2$ or $\mathbb{S}^2$

Model	α_t	α_s	α_e	Type	ν_t	ν_s
A: DEMF(1,0,2)	1	0	2	Separable order 1	$1/2$	1
B: DEMF(1,2,1)	1	2	1	Critical diffusion	$1/2$	1
C: DEMF(2,0,2)	2	0	2	Separable order 2	$3/2$	1
D: DEMF(2,2,0)	2	2	0	Iterated diffusion	1	2

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

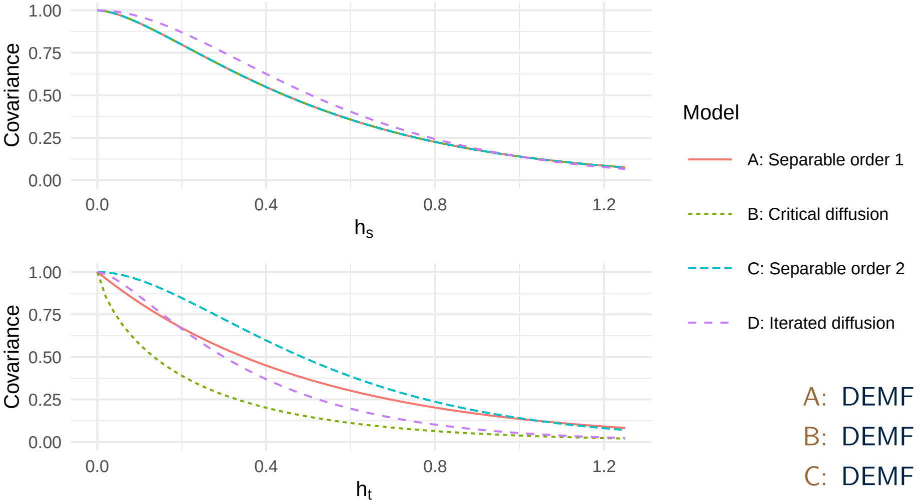

A: DEMF(1,0,1)
B: DEMF(1,2,1)
C: DEMF(2,0,1)
D: DEMF(2,2,0)

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

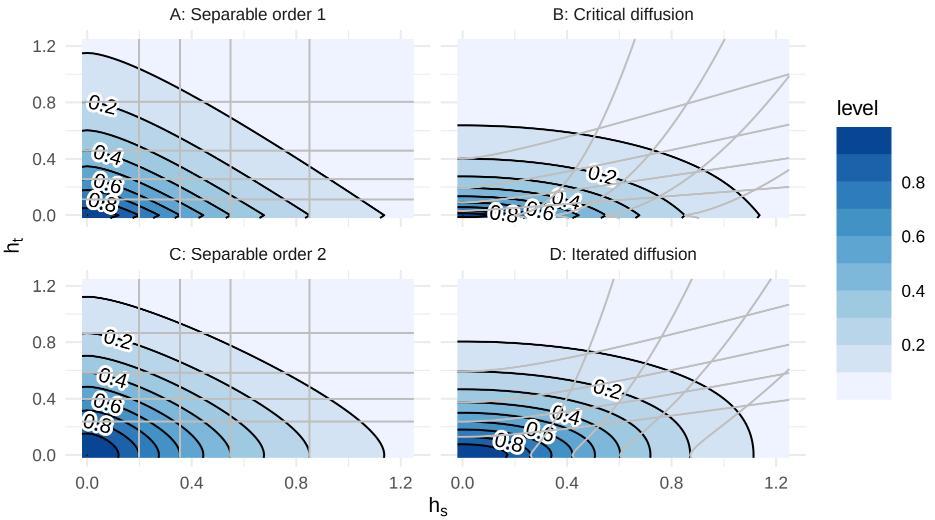

Prediction

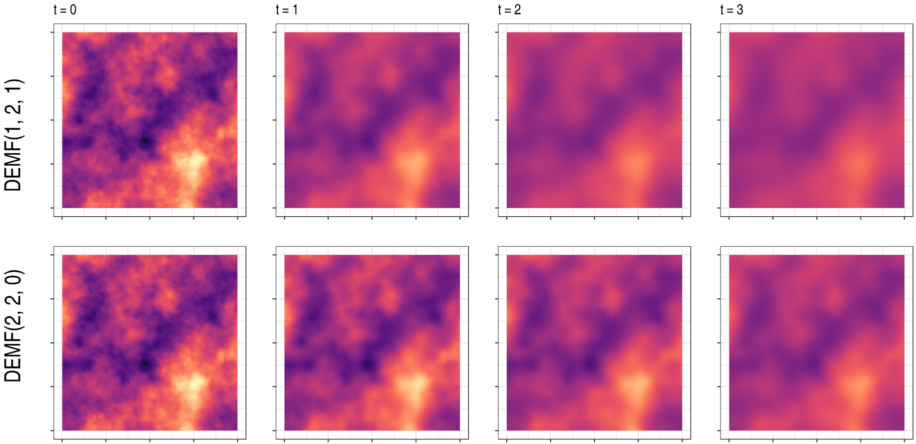

Conditional expectations into the future diffuse across space; some spatial dynamics.

Example: Global temperature modelling

$$
\begin{aligned}
\eta(\mathbf{s}, t) &= \mu + \alpha E(\mathbf{s}) + b(\mathbf{s}, t) + v(\mathbf{s}, t) + u(\mathbf{s}, t) \\
y_i &= \eta(\mathbf{s}_i, t_i) + \epsilon_i
\end{aligned}
$$

- μ: global offset
- $E(\mathbf{s})$: elevation
- $b(\mathbf{s}, t)$: seasonal effect varying by latitude and time
- $v(\mathbf{s}, t)$: slowly varying spatial effect
- $u(\mathbf{s}, t)$: daily weather component, using a DEMF model

Time series and computational mesh

Daily temperature at 13581 stations for year 2022.

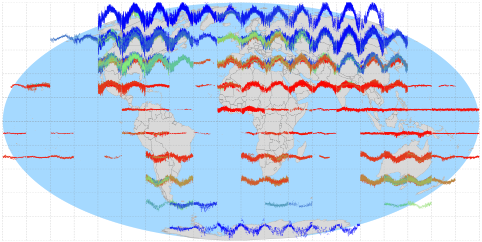

Spatial mesh with 1251 nodes, and the 13581 stations.

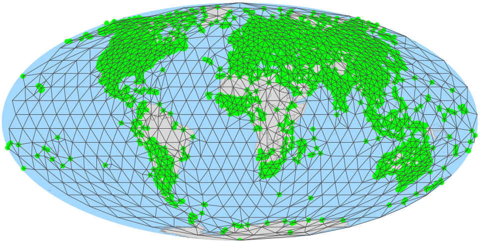

Seasonal latitude effect

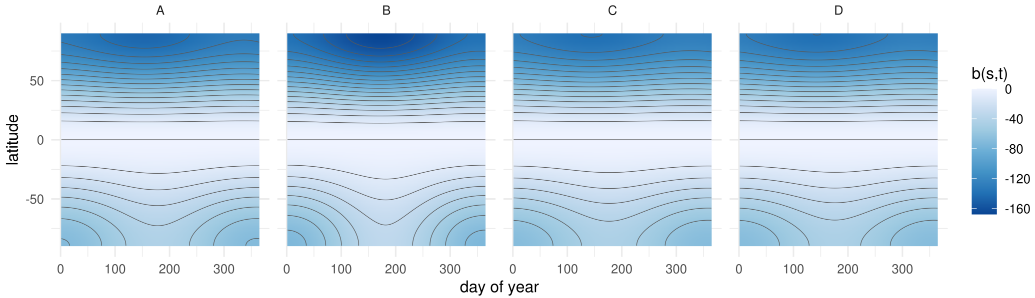

Daily weather component for days 46, 137, 228, and 319 of 2022

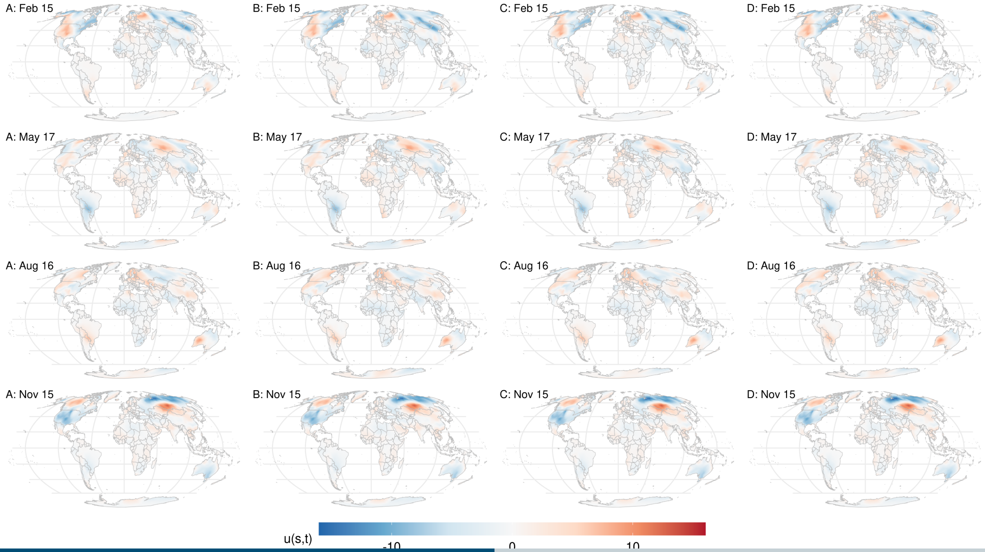

Local model fit and prediction scores

Model	No u	A	B	C	D
R^2	0.8649	0.9718	**0.9718**	0.9718	0.9718
DIC	5.8103	4.3217	4.3222	**4.3212**	4.3222
WAIC	5.8101	4.3123	4.3128	**4.3118**	4.3128
LPO	2.9046	2.1332	**2.1330**	2.1335	2.1331
LCPO	2.9051	2.1572	2.1574	**2.1569**	2.1574
MSE	19.5151	4.1163	**4.1136**	4.1193	4.1155
MAE	3.3236	1.4667	**1.4657**	1.4678	1.4661
CRPS	2.4288	1.0947	**1.0943**	1.0952	1.0945
SCRPS	1.7905	1.3947	**1.3945**	1.3949	1.3946

Multi-day forecast errors

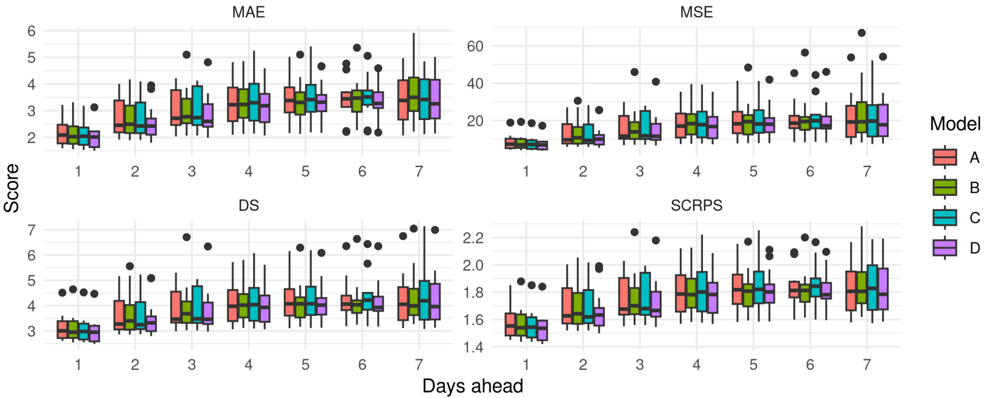

Extensions to metric graphs

- Let Γ be a metric graph, defined in terms of a set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E}$ connecting the vertices.
- The difference to a regular graph is that the edges are defined as rectifiable curves, and a position $s \in \Gamma$ can be represented as (e, t), where $e \in \mathcal{E}$ denotes an edge, and t is a position on that edge. Thus, these spaces contain linear networks as a special case.
- Whittle–Matérn fields on metric graphs (Bolin et al 2023, Bernoulli) as the solution to

$$(\kappa^2 - \Delta_\Gamma)^{\alpha/2}(\tau u) = \mathrm{d}\mathcal{W}, \quad \text{on } \Gamma$$

where $\alpha > 1/2$, $\mathrm{d}\mathcal{W}$ is Gaussian white noise on Γ and Δ_Γ is the so-called Kirchhoff Laplacian on Γ, which is an operator that acts as the second derivative on the edges.
- Non-separable space-time models can be defined on metric graphs using the same technique as in our current paper:

$$\left(\gamma_t \frac{d}{dt} + (\kappa^2 + \rho\,\mathrm{d}_s - \Delta_\Gamma)^{\alpha_s/2}\right)^{\alpha_t} u = \mathrm{d}\mathcal{E}_\mathcal{Q}, \quad \text{on } \Gamma \times [0, T],$$

where d_s is a transport/advection term
- These spatial models are implemented in the R package `MetricGraph`, which also contains an implementation of LGCPs on metric graphs.

Extensions to metric graphs

- Let Γ be a metric graph, defined in terms of a set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E}$ connecting the vertices.
- The difference to a regular graph is that the edges are defined as rectifiable curves, and a position $s \in \Gamma$ can be represented as (e, t), where $e \in \mathcal{E}$ denotes an edge, and t is a position on that edge. Thus, these spaces contain linear networks as a special case.
- Whittle–Matérn fields on metric graphs (Bolin et al 2023, Bernoulli) as the solution to

$$(\kappa^2 - \Delta_\Gamma)^{\alpha/2}(\tau u) = \mathrm{d}\mathcal{W}, \quad \text{on } \Gamma$$

  where $\alpha > 1/2$, $\mathrm{d}\mathcal{W}$ is Gaussian white noise on Γ and Δ_Γ is the so-called Kirchhoff Laplacian on Γ, which is an operator that acts as the second derivative on the edges.
- Non-separable space-time models can be defined on metric graphs using the same technique as in our current paper:

$$(\gamma_t \frac{d}{dt} + (\kappa^2 + \rho\,\mathrm{d}_s - \Delta_\Gamma)^{\alpha_s/2})^{\alpha_t} u = \mathrm{d}\mathcal{E}_{\mathcal{Q}}, \quad \text{on } \Gamma \times [0, T],$$

  where d_s is a transport/advection term
- These spatial models are implemented in the R package `MetricGraph`, which also contains an implementation of LGCPs on metric graphs.

Example: LGCP on a metric graph

Simulated LGCP on a metric graph and estimate of log-intensity based on R-INLA and the MetricGraph package.

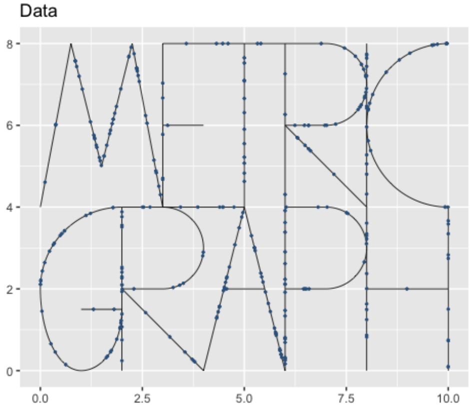

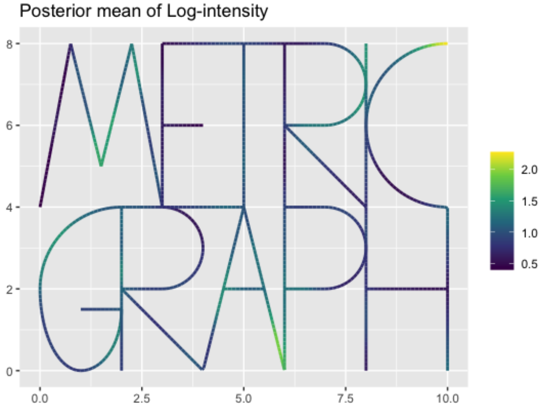

Summary of separability concepts

Basics:

- Functions, covariance, precision
- Additive, multiplicative, additive multiplication combinations
- Non-separability needed for realistic dynamics

Further concepts

- Anisotropy
- Non-stationarity; separable and non-separable, space and/or time
- Asymmetry; transport/advection terms in the space-time operator
- Manifold domains; easy in practice and most of the theory; some theory more difficult

Practical considerations

- Non-separable models are more expensive to fit
- Interpolation/infilling applications show little difference from separable models
- Full temporal predictions are fundamentally different;
  non-separability can be essential

Summary

- Non-separable space-time models result from simple stochastic PDEs
- Model assessment should target the intended use case, e.g. interpolation *or* prediction
- 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
- INLAspacetime implementation package:
  https://github.com/eliaskrainski/INLAspacetime
- R-INLA documentation and examples:
  https://www.r-inla.org/
- inlabru/fmesher documentation and examples:
  https://inlabru-org.github.io/inlabru/ (and .../fmesher/)
- MetricGraph documentation and examples:
  https://davidbolin.github.io/MetricGraph/

Summary

- Non-separable space-time models result from simple stochastic PDEs
- Model assessment should target the intended use case, e.g. interpolation *or* prediction
- 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
- INLAspacetime implementation package:
  https://github.com/eliaskrainski/INLAspacetime
- R-INLA documentation and examples:
  https://www.r-inla.org/
- inlabru/fmesher documentation and examples:
  https://inlabru-org.github.io/inlabru/ (and .../fmesher/)
- MetricGraph documentation and examples:
  https://davidbolin.github.io/MetricGraph/

Summary

- Non-separable space-time models result from simple stochastic PDEs
- Model assessment should target the intended use case, e.g. interpolation *or* prediction
- 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`
- `INLAspacetime` implementation package:
  `https://github.com/eliaskrainski/INLAspacetime`
- `R-INLA` documentation and examples:
  `https://www.r-inla.org/`
- `inlabru/fmesher` documentation and examples:
  `https://inlabru-org.github.io/inlabru/` (and `.../fmesher/`)
- `MetricGraph` documentation and examples:
  `https://davidbolin.github.io/MetricGraph/`