

Embedding stochastic PDEs in Bayesian spatial statistics software

Finn Lindgren

finn.lindgren@ed.ac.uk



NASPDE, EURANDOM, Eindhoven, 17 May 2023

GAMs and general kriging

- Linear GAMs with GPs on space and covariates:

$$\eta_i = \sum_k v_k(z_{ik}) + u(\mathbf{s}_i),$$

each $v_k(\cdot)$ and $u(\cdot)$ represented with basis expansions with jointly Gaussian coefficients \mathbf{x} .

- Linear observations with additive Gaussian observation noise: $\mathbf{y} = \boldsymbol{\eta} + \boldsymbol{\epsilon} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}$
- Covariance kriging

$$\boldsymbol{\Sigma}_y = \mathbf{A}\boldsymbol{\Sigma}_x\mathbf{A}^\top + \boldsymbol{\Sigma}_\epsilon$$

$$E(\mathbf{x}|\mathbf{y}) = \boldsymbol{\mu} + \boldsymbol{\Sigma}_x\mathbf{A}^\top\boldsymbol{\Sigma}_y^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu})$$

- Precision kriging

$$\mathbf{Q}_{x|y} = \mathbf{Q}_x + \mathbf{A}^\top\mathbf{Q}_\epsilon\mathbf{A}$$

$$E(\mathbf{x}|\mathbf{y}) = \boldsymbol{\mu} + \mathbf{Q}_{x|y}^{-1}\mathbf{A}^\top\mathbf{Q}_\epsilon(\mathbf{y} - \mathbf{A}\boldsymbol{\mu})$$

- Non-Gaussian observations with link function: $E(y_i|\boldsymbol{\theta}, \mathbf{x}) = h(\eta_i)$

Observation level covariance vs latent level precision

- Covariance kriging: linear solve with a Σ , $\Sigma_{ij} = \text{Cov}(y_i, y_j)$
- Precision kriging: linear solve with a Q , $Q_{ij} = \text{Prec}(x_i, x_j | \mathbf{y})$
 $Q = LL^\top$ for a given latent variable ordering, and sparse lower triangular L with the sparsity from Q plus Cholesky infill.
The prior Q_x for GRF/SPDE process components are obtained via a local Finite Element construction, giving the model in a chosen finite function space closest to the full model.

Finite element structure

Matérn-Whittle processes

Linear Gaussian process/field representations via SPDEs:

$$(\kappa^2 - \Delta)^\alpha u(\mathbf{s}) \, d\mathbf{s} = d\mathcal{W}(\mathbf{s}) \kappa^{\alpha-d/2} / \tau$$

For constant parameters, $u(\mathbf{s})$ has spatial Matérn covariance on \mathbb{R}^d , and generalised Matérn-Whittle covariance on general manifolds. The smoothness index is $\nu = \alpha - d/2$ and the variance is proportional to $1/\tau^2$. Whittle (1954, 1963), Lindgren et al (2011)

Discrete domain Gaussian Markov random fields (GMRFs)

$\mathbf{x} = (x_1, \dots, x_n) \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$ is Markov with respect to a neighbourhood structure $\{\mathcal{N}_i, i = 1, \dots, n\}$ if $Q_{ij} = 0$ whenever $j \notin \mathcal{N}_i \cup i$.

- Continuous domain basis representation with weights: $x(\mathbf{s}) = \sum_{k=1}^n \psi_k(\mathbf{s}) u_k$
- Project the SPDE solution space onto local basis functions:
random Markov dependent basis weights (Lindgren et al, 2011).

Non-stationarity

Non-stationary Matérn-Whittle processes

The Sampson & Guttorp (1992) deformation method motivates a non-stationary generalisation on \mathbb{R}^2 :

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

where $\kappa(\mathbf{s})$ and $\mathbf{H}(\mathbf{s})$ are derived from the metric tensor of the deformation. For deformation *not* from \mathbb{R}^d onto \mathbb{R}^d , this non-stationary model is distinct from the deformation method, but keeps much of the intuition, as the variance will be approximatively independent of $\kappa(\mathbf{s})$.

RKHS inner products of linear SPDEs

The spatial solutions $u(\mathbf{s})$ to

$$\mathcal{L}u(\mathbf{s}) d\mathbf{s} = d\mathcal{W}(\mathbf{s}) \quad \text{where } d\mathcal{W}(\mathbf{s}) \text{ is white noise on } \Omega$$

have RKHS inner product

$$Q_\Omega(f, g) = \langle \mathcal{L}f, \mathcal{L}g \rangle_\Omega$$

plus potential boundary terms.

Non-separable space-time: Matérn driven heat equation on the sphere

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

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

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:

$$\nu_t = \min \left[\alpha_t - \frac{1}{2}, \frac{\nu_s}{\alpha_s} \right], \quad \alpha_t = \nu_t \max \left(1, \frac{\beta_s}{\beta_*(\nu_s, d)} \right) + \frac{1}{2},$$

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

$$\beta_*(\nu_s, d) = \frac{\nu_s}{\nu_s + d/2}, \quad \alpha_e = \frac{1 - \beta_s}{\beta_*(\nu_s, d)} \nu_s = (\nu_s + d/2)(1 - \beta_s).$$

Spectra and finite element structure

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

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

- 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

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

- The finite element approximation has structure

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

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

Latent Gaussian models

Hierarchical model with latent jointly Gaussian variables

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}) \quad (\text{covariance parameters})$$

$$(\mathbf{u} \mid \boldsymbol{\theta}) \sim \text{N}(\boldsymbol{\mu}_u, \mathbf{Q}_u^{-1}) \quad (\text{latent Gaussian variables})$$

$$(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta}) \sim p(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta}) \quad (\text{observation model})$$

We are interested in the posterior densities $p(\boldsymbol{\theta} \mid \mathbf{y})$, $p(\mathbf{u} \mid \mathbf{y})$ and $p(u_i \mid \mathbf{y})$.

Approximate conditional posterior distribution

Let $\hat{\mathbf{u}}(\boldsymbol{\theta})$ be the mode of the posterior density $p(\mathbf{u} \mid \mathbf{y}, \boldsymbol{\theta}) \propto p(\mathbf{u} \mid \boldsymbol{\theta})p(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta})$. Construct an approximate conditional posterior distribution, via Newton optimisation for \mathbf{u} given $\boldsymbol{\theta}$:

$$p_G(\mathbf{u} \mid \mathbf{y}, \boldsymbol{\theta}) \sim \text{N}(\hat{\boldsymbol{\mu}}, \hat{\mathbf{Q}}^{-1})$$

$$\mathbf{0} = \nabla_{\mathbf{u}} \{ \ln p(\mathbf{u} \mid \boldsymbol{\theta}) + \ln p(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta}) \} \Big|_{\mathbf{u}=\hat{\boldsymbol{\mu}}(\boldsymbol{\theta})}$$

$$\hat{\mathbf{Q}} = \mathbf{Q}_u - \nabla_{\mathbf{u}}^2 \ln p(\mathbf{y} \mid \mathbf{u}, \boldsymbol{\theta}) \Big|_{\mathbf{u}=\hat{\boldsymbol{\mu}}(\boldsymbol{\theta})}$$

Classic and compact INLA methods (\sim description)

- Laplace approximation at the conditional posterior mode \mathbf{x}^* , and uncertainty integration:

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto \frac{p(\boldsymbol{\theta})p(\mathbf{x}|\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x})}{p(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})} \Bigg|_{\mathbf{x}=\mathbf{x}^*} \approx \frac{p(\boldsymbol{\theta})p(\mathbf{x}|\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x})}{p_G(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})} \Bigg|_{\mathbf{x}=\mathbf{x}^*} = \hat{p}(\boldsymbol{\theta}|\mathbf{y})$$

$$p(x_i|\mathbf{y}) = \int p(x_i|\boldsymbol{\theta}, \mathbf{y})p(\boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta} \approx \sum_k \hat{p}(x_i|\boldsymbol{\theta}^{(k)}, \mathbf{y})\hat{p}(\boldsymbol{\theta}^{(k)}|\mathbf{y})w_k = \hat{p}(x_i|\mathbf{y})$$

- Let $\hat{\boldsymbol{\mu}} = \mathbb{E}(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$ and $\mathbf{Q}_\epsilon = -\nabla_{\mathbf{x}} \nabla_{\mathbf{x}}^\top \log p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}^*)$

- Classic method: Laplace approximation of each $\hat{p}(x_i|\boldsymbol{\theta}, \mathbf{y})$, and

$$\left\{ \begin{bmatrix} \mathbf{Ax} \\ \mathbf{x} \end{bmatrix} \middle| \boldsymbol{\theta}, \mathbf{y} \right\} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{A}\hat{\boldsymbol{\mu}} \\ \hat{\boldsymbol{\mu}} \end{bmatrix}, \begin{bmatrix} \mathbf{Q}_\epsilon + \delta \mathbf{I} & -\delta \mathbf{A} \\ -\delta \mathbf{A}^\top & \mathbf{Q}_x + \delta \mathbf{A}^\top \mathbf{A} \end{bmatrix}^{-1} \right), \text{ with } \delta \gg 0$$

- Compact method: Variational approximation of $\hat{p}(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$, and $\{\mathbf{x}|\boldsymbol{\theta}, \mathbf{y}\} \sim \mathcal{N}(\hat{\boldsymbol{\mu}}, [\mathbf{Q}_x + \mathbf{A}^\top \mathbf{Q}_\epsilon \mathbf{A}]^{-1})$

inlabru software interface concepts

- Model components are declared similarly to R-INLA:

```
# INLA:
~ covar + f(name, model = ...)
# inlabru
~ covar + name(input, model = ...)
~ covar # is translated into...
~ covar(covar, model = "linear")
~ name(1) # Used for intercept-like components
```

- In R-INLA, $\boldsymbol{\eta} = \mathbf{A}\mathbf{u} = \mathbf{A}_0 \sum_{k=1}^K \mathbf{A}_k \mathbf{u}_k$, where the rows of \mathbf{A}_k only extract individual elements from each \mathbf{u}_k , and the overall \mathbf{A}_0 is user defined (via `inla.stack()`).
- In inlabru, $\boldsymbol{\eta} = h(\mathbf{u}_1, \dots, \mathbf{u}_K, \mathbf{A}_1 \mathbf{u}_1, \dots, \mathbf{A}_K \mathbf{u}_K)$, where $h(\cdot)$ is a general R expression of named latent components \mathbf{u}_k and intermediate "effects" $\mathbf{A}_k \mathbf{u}_k$
- \mathbf{A}_k by default acts either as in R-INLA, or is determined by a *mapper* method. Predefined default mappers include e.g. spatial evaluation of SPDE/GRMF models that map between coordinates and meshes, and mappers that combine other mappers (used to combine main/group/replicate for all components)

Input mappers

- Each named component has main/group/replicate *inputs*, that are given to the mappers to evaluate A_k . For a given latent *state*, the resulting *effect* values are made available to the predictor expression.

```
bru_mapper() # generic
bru_mapper_index(n) # Basic index mapping
bru_mapper_linear() # Basic linear mapping
bru_mapper_matrix(labels) # Basic linear multivariable mapping
bru_mapper_factor(values, factor_mapping) # Factor variable mapping
bru_mapper_multi(mappers) # kronecker product components
bru_mapper_collect(mappers, hidden) # For concatenated components, like bym
bru_mapper_const() # Constants
bru_mapper.inla.mesh(mesh) # 2D and spherical mesh mappings
bru_mapper.inla.mesh.1d(mesh) # Interval and cyclic interval mappings
```

- Common methods that return essential characteristics

```
ibm_n(mapper) # The size of the latent component
ibm_values(mapper) # The covariate/index "values" given to INLA
ibm_jacobian(mapper, input) # The "A-matrix" for given input values
```

- Model component definition example:

```
comp <- ~ -1 + field(cbind(easting, northing), model = spde) + param(1)
```

- Predictor formula examples, including naming of the response variable:

```
form1 <- my_counts ~ param + field
form2 <- response ~ exp(param) + exp(field)
```

- Main method call structure:

```
bru(components = comp,
     like(formula = form1, family = "poisson", data = data1),
     like(formula = form2, family = "normal", data = data2))
```

- Simplified notations for common special cases;

```
formula = response ~ .
```

gives the full additive model of all the available components, or

```
components = response ~ Intercept(1) + field(...
```

Plain INLA code for space-time model

```
matern <- inla.spde2.pcmatern(mesh, ...)  
  
field_A <- inla.spde.make.A(mesh,  
                           coordinates(data),  
                           group = data$year - min(data$year) + 1,  
                           n.group = 10)  
stk <- inla.stack(data = list(response = data$response),  
                 A = list(field_A, 1),  
                 effects = list(field_index, list(covar = data$covar)))  
  
formula <- response ~ 1 + covar +  
  f(field, model = matern, group = field_group, control.group = ...)  
  
fit <- inla(formula = formula,  
           data = inla.stack.data(stk, matern = matern),  
           family = "normal",  
           control.predictor = list(A = inla.stack.A(stk)))
```

inlabru code for space-time model

```
matern <- inla.spde2.pcmatern(mesh, ...)  
  
year_mapper <- bru_mapper(inla.mesh.1d(sort(unique(data$year))), indexed = TRUE)  
  
comp <- response ~ Intercept(1) + covar +  
  field(coordinates, model = matern, group = year, group_mapper = year_mapper,  
    control.group = ...)  
  
fit <- bru(components = comp,  
  data = data,  
  family = "normal")
```

Implied:

- `coordinates` \rightarrow `sp::coordinates(.data.)`
- `formula = response ~ .`
- `data` and `family` passed on to a `like()` call

Latent Gaussian models of R-INLA type and inlabru extension

LGM of R-INLA type

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}), \quad (\text{hyper-})\text{parameters}$$

$$\mathbf{u}|\boldsymbol{\theta} \sim \text{N}(\boldsymbol{\mu}_u, \mathbf{Q}(\boldsymbol{\theta})^{-1}), \quad \text{complex structured latent Gaussian field}$$

$$\boldsymbol{\eta}(\mathbf{u}) = \mathbf{A}\mathbf{u}, \quad \text{linear predictor, linear combination of the latent variables}$$

$$y_i|\mathbf{u}, \boldsymbol{\theta} \sim p(y_i|\eta_i(\mathbf{u}), \boldsymbol{\theta}), \quad \text{response variables } y_i, \text{ conditionally independent}$$

Extended LGM of inlabru type

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}), \quad (\text{hyper-})\text{parameters}$$

$$\mathbf{u}|\boldsymbol{\theta} \sim \text{N}(\boldsymbol{\mu}_u, \mathbf{Q}(\boldsymbol{\theta})^{-1}), \quad \text{complex structured latent Gaussian field}$$

$$\boldsymbol{\eta}(\mathbf{u}) = h(\mathbf{u}), \quad \text{non-linear predictor, general function of the latent variables}$$

$$y_i|\mathbf{u}, \boldsymbol{\theta} \sim p(y_i|\eta_i(\mathbf{u}), \boldsymbol{\theta}), \quad \text{response variables } y_i, \text{ conditionally independent}$$

Approximate INLA for non-linear predictors

Linearised predictor

Let $\tilde{\eta}(\mathbf{u})$ be the non-linear predictor, and let $\bar{\eta}(\mathbf{u})$ be the 1st order Taylor approximation at some \mathbf{u}_0 ,

$$\bar{\eta}(\mathbf{u}) = \tilde{\eta}(\mathbf{u}_0) + \mathbf{B}(\mathbf{u} - \mathbf{u}_0) = [\tilde{\eta}(\mathbf{u}_0) - \mathbf{B}\mathbf{u}_0] + \mathbf{B}\mathbf{u},$$

where \mathbf{B} is the derivative matrix for the non-linear predictor, evaluated at \mathbf{u}_0 .

The non-linear observation model $\tilde{p}(\mathbf{y}|\mathbf{u}, \boldsymbol{\theta})$ is approximated by

$$\bar{p}(\mathbf{y}|\mathbf{u}, \boldsymbol{\theta}) = p(\mathbf{y}|\bar{\eta}(\mathbf{u}), \boldsymbol{\theta}) \approx p(\mathbf{y}|\tilde{\eta}(\mathbf{u}), \boldsymbol{\theta}) = \tilde{p}(\mathbf{y}|\mathbf{u}, \boldsymbol{\theta})$$

The non-linear model posterior is factorised as

$$\tilde{p}(\boldsymbol{\theta}, \mathbf{u}|\mathbf{y}) = \tilde{p}(\boldsymbol{\theta}|\mathbf{y})\tilde{p}(\mathbf{u}|\mathbf{y}, \boldsymbol{\theta}),$$

and the linear model approximation is factorised as

$$\bar{p}(\boldsymbol{\theta}, \mathbf{u}|\mathbf{y}) = \bar{p}(\boldsymbol{\theta}|\mathbf{y})\bar{p}(\mathbf{u}|\mathbf{y}, \boldsymbol{\theta}).$$

Iterated INLA in inlabru

The observation model is linked to \mathbf{u} only through the non-linear predictor $\tilde{\eta}(\mathbf{u})$.

Iterative INLA algorithm:


- 1 Let \mathbf{u}_0 be an initial linearisation point for the latent variables.
- 2 Compute the predictor linearisation at \mathbf{u}_0
- 3 Compute the linearised INLA posterior $\bar{p}(\boldsymbol{\theta}|\mathbf{y})$ and let $\hat{\boldsymbol{\theta}} = \operatorname{argmax}_{\boldsymbol{\theta}} \bar{p}(\boldsymbol{\theta}|\mathbf{y})$
- 4 Let $\mathbf{u}_1 = \operatorname{argmax}_{\mathbf{u}} \bar{p}(\mathbf{u}|\mathbf{y}, \hat{\boldsymbol{\theta}})$ be the initial candidate for new linearisation point.
- 5 Let $\mathbf{u}_{\alpha} = (1 - \alpha)\mathbf{u}_0 + \alpha\mathbf{u}_1$, and find the value α that minimises $\|\tilde{\eta}(\mathbf{u}_{\alpha}) - \bar{\eta}(\mathbf{u}_1)\|$.
- 6 Set the linearisation point \mathbf{u}_0 to \mathbf{u}_{α} and repeat from step 2, unless the iteration has converged to a given tolerance.
- 7 Compute $\bar{p}(\mathbf{u}|\mathbf{y})$

In step 4, only the *conditional* posterior mode for \mathbf{u} is needed, so the costly nested integration step of the R-INLA algorithm only needs to be run in a final iteration of the algorithm, in step 7.

Step 5 can use an approximate line search method.

Example: Thinned Poisson point processes

We want to model the presence of groups of dolphins using a Log-Gaussian Cox Process (LGCP)
However, when surveying dolphins from a ship travelling along lines (*transects*), the probability of detecting a group of animals depends their distance distance from the ship.



Example: Thinned Poisson point processes

We want to model the presence of groups of dolphins using a Log-Gaussian Cox Process (LGCP). However, when surveying dolphins from a ship travelling along lines (*transects*), the probability of detecting a group of animals depends their distance distance from the ship, e.g. via

$$P(\text{detection}) = 1 - \exp\left(-\frac{\sigma}{\text{distance}}\right) \quad (\text{hazard rate model})$$

This results in a *thinned* Poisson process model on (space, distance) along the transects:

$$\log(\lambda(\mathbf{s}, \text{distance})) = \text{Intercept} + \text{field}(\mathbf{s}) + \log [P(\text{detection at } \mathbf{s} \mid \text{distance}, \sigma)] + \log(2)$$

inlabru knows how to construct the Poisson process likelihood along lines and on polygons, and kronecker spaces (line \times distance)

We can define $\log(\sigma)$ as a latent Gaussian variable and iteratively linearise. The non-linearity is mild, and the iterative INLA method converges.


```
log_det_prob <- function(distance, log_sig) {  
  log1p(-exp(-exp(log_sig) / distance))  
}  
  
comp <- ~ field(coordinates, model = matern) + log_sig(1) + Intercept(1)  
form <- coordinates + distance ~  
  Intercept + field + log_det_prob(distance, log_sig) + log(2)  
  
fit <- bru(  
  components = comp,  
  like(  
    family = "cp", formula = form,  
    data = mexdolphins$points, # sp::SpatialPointsDataFrame  
    samplers = mexdolphins$samplers, # sp::SpatialLinesDataFrame  
    domain = list(  
      coordinates = mexdolphins$mesh,  
      distance = INLA::inla.mesh.1d(seq(0, 8, length.out = 30))  
    )  
  )  
)  
)
```

Posterior prediction method

```
pred_points <- pixels(mexdolphins$mesh, nx = 200, ny = 100, mask = mexdolphins$ppoly)
pred <- predict(fit, pred_points, ~ exp(field + Intercept))
```

```
det_prob <- function(distance, log_sig) { 1 - exp(-exp(log_sig) / distance) }
pred_dist <- data.frame(distance = seq(0, 8, length = 100))
det_prob <- predict(fit, pred_dist, ~ det_prob(distance, log_sig), include = "log_sig")
```

```
ggplot() + gg(pred) + gg(mexdolphins$samplers) + gg(mexdolphins$ppoly) + ...
```



Data level prediction

47 groups were seen. How many would be seen along the transects under perfect detection?

```
predpts_transect <- ipoints(mexdolphinsamplers, mexdolphinsmesh)
Lambda_transect <- predict(fit, predpts_transect,
  ~ 16 * sum(weight * exp(field + Intercept)))
```

mean	sd	q0.025	q0.5	q0.975	median	mean.mc_std_err	sd.mc_std_err
85.93922	24.93825	42.0879	82.61408	149.988	82.61408	2.493825	2.239346


How many would be seen under perfect detection across the whole study area?

```
predpts <- ipoints(mexdolphinsppoly, mexdolphinsmesh)
Lambda <- predict(fit, predpts, ~ sum(weight * exp(field + Intercept)))
```


mean	sd	q0.025	q0.5	q0.975	median	mean.mc_std_err	sd.mc_std_err
319.752	81.54696	190.0406	313.1301	520.4723	313.1301	8.154696	8.796282

Integration points


Line integration



Area integration



Computational mesh




Complex prediction expressions

What's the predictive distribution of group counts?

```
Ns <- seq(50, 650, by = 1)
Nest <- predict(
  fit,
  predpts,
  ~ data.frame(
    N = Ns,
    density = dpois(Ns, lambda = sum(weight * exp(field + Intercept)))
  ),
  n.samples = 2500
)

Nest$plugin_estimate <- dpois(Nest$N, lambda = Lambda$mean)
```


Full posterior prediction uncertainty vs plugin prediction



EUSTACE ANALYSIS

Combines in-situ and satellite data sources to derive daily air temperatures across the globe with quantified uncertainties.


- Daily mean air temperature (2 m) estimates from the mid-late 19th century at $\frac{1}{4}$ degree resolution.
- Observational dataset for use in climate monitoring, services and research.
 - Quantify bias and uncertainty arising from observational sampling (in space and time);
 - Quantify uncertainty from instrumental effects/network changes.
- Higher resolution daily gridded analyses for regional climate
 - Combine in situ and remote sensing data to support high resolution analysis.
 - Absolute temperature rather than anomaly product.




ENSEMBLE ANALYSIS

- Samples drawn from joint posterior distribution of temperature and bias variables.
- Temperature model samples projected onto analysis grid.
- Spatial/temporal correlation in analysis errors is encoded into the ensemble.
- Summary statistics can be derived from the ensemble. Expected value, total uncertainty and observation constraint information also available.

EUSTACE Ensemble 04/08/2003-13/08/2003




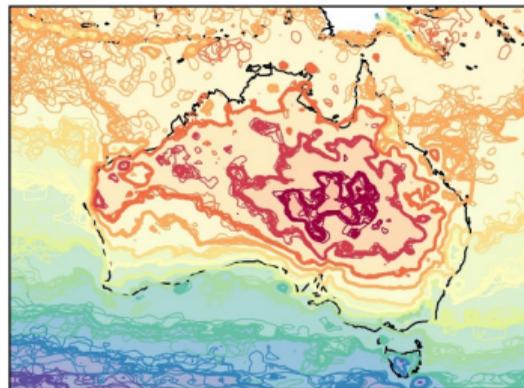
EUSTACE Ensemble 30/07/2010-05/08/2010



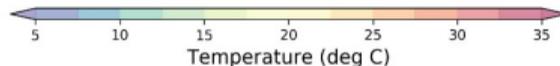
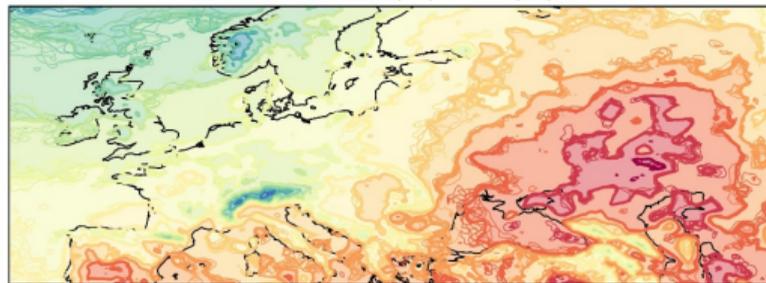
ENSEMBLE ANALYSIS

- Samples drawn from joint posterior distribution of temperature and bias variables.
- Temperature model samples projected onto analysis grid.
- Spatial/temporal correlation in analysis errors is encoded into the ensemble.
- Summary statistics can be derived from the ensemble. Expected value, total uncertainty and observation constraint information also available.

EUSTACE Ensemble 01/01/2006-14/01/2006



EUSTACE Ensemble 30/07/2010-05/08/2010





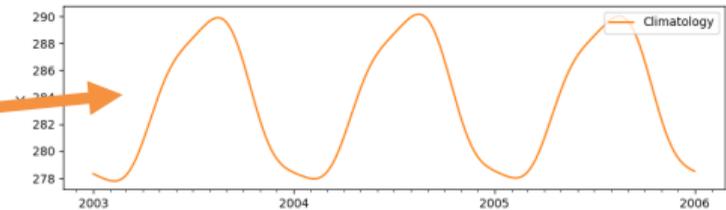
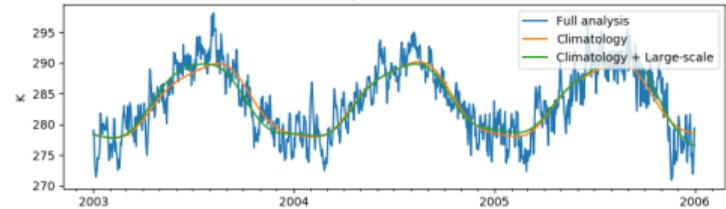
MULTI-SCALE ANALYSIS MODEL

Statistical model for temperature variations and different scales (space and time):

- **Climatological variation**: local seasonal cycle with effects of latitude, altitude and coastal influence.
- **Large-scale variation**: Slowly varying climatological mean temperature field.
- **Daily Local**: daily variability associated with weather.

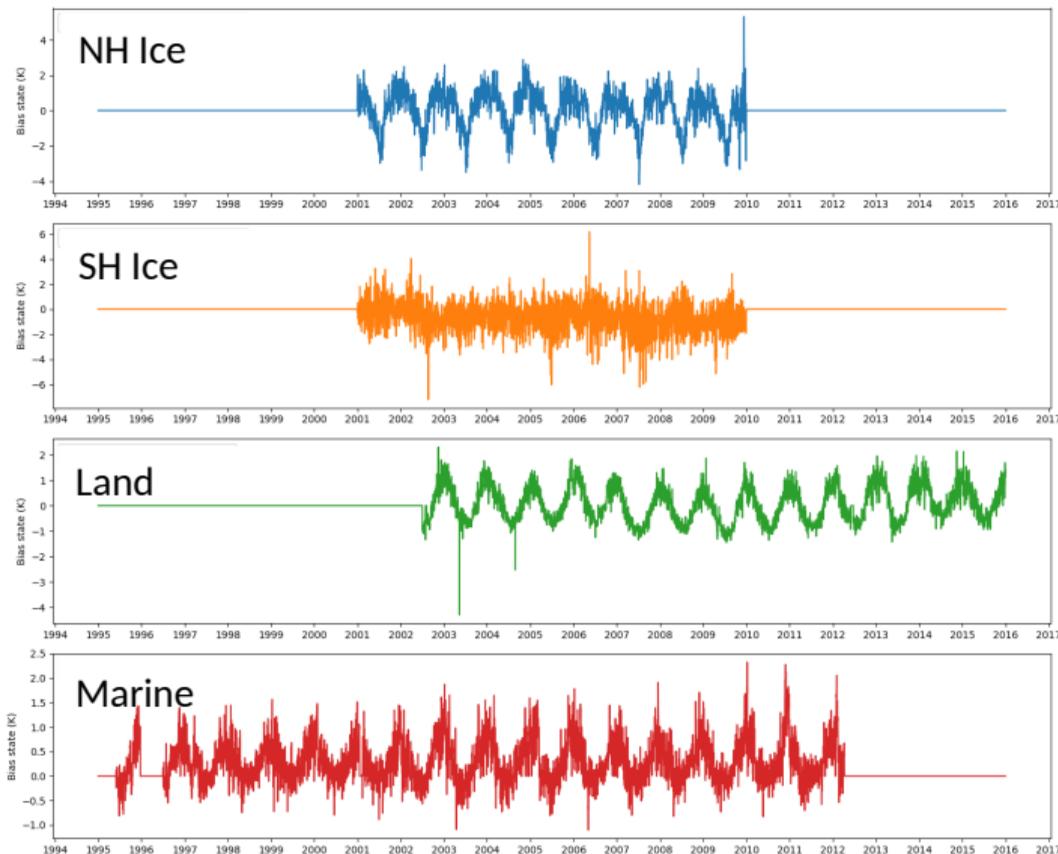
Simultaneously estimates observational biases of known bias structures:

- e.g. satellite biases, station homogenisation.

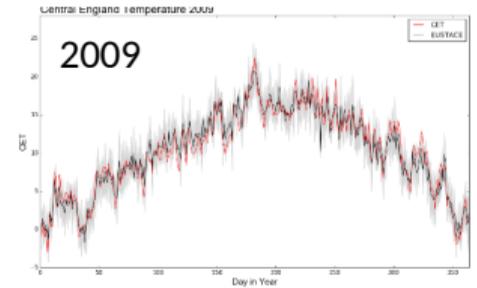
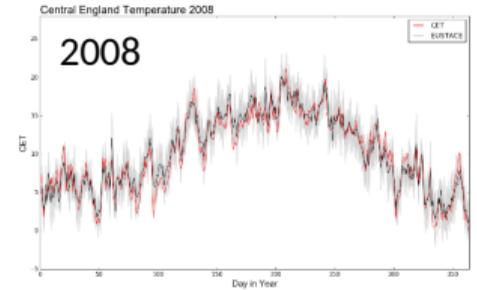
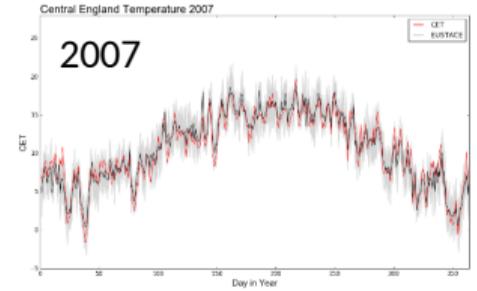
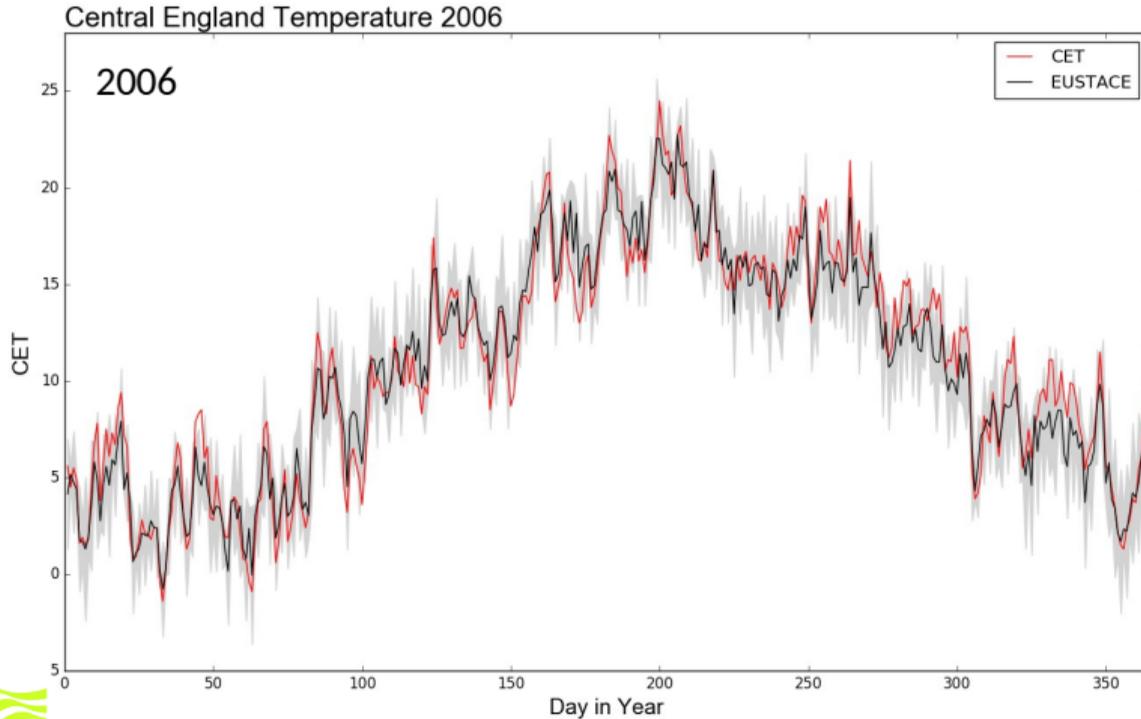


SATELLITE BIAS MODELS

- Simplified model of known error structures in satellite air temperature retrievals:
 - Global/hemispheric systematic bias covariates.
 - Daily estimates of spatially varying bias as a spatial random field.
- Estimated jointly with daily temperature variability.



COMPARING EUSTACE WITH CENTRAL ENGLAND TEMPERATURE



MULTI-SCALE ANALYSIS MODEL

Statistical model for temperature variations and different scales (space and time):

- **Climatological variation:** local seasonal cycle with effects of latitude, altitude and coastal influence.
- **Large-scale variation:** Slowly varying climatological mean temperature field. Station homogenisation.
- **Daily Local:** daily variability associated with weather. Satellite retrieval biases.

Simultaneously estimates observational biases of known bias structures:

- e.g. satellite biases, station homogenisation.

Processed on STFC's LOTUS cluster www.jasmin.ac.uk:

- Largest solves processed on 20 core/256GB RAM node.
- Highly parallel observation pre-processing.

Element	Resolution	N Variables
Seasonal	Bimonthly x 1° SPDE	245,772
Slow-scale*	5 year x 5° SPDE	107,604
Latitude	0.5° latitude SPDE	721
Altitude	(0.25° grid)	1
Coastal	(0.25° grid)	1
Grand mean	Analysis mean	1

Element	Resolution	N Variables
Large-scale	3 monthly x 5° SPDE	1,752,408
Station bias	NA	82,072

Element	Resolution	N Variables per day
Daily local	~0.5 degree SPDE	162,842
Satellite bias (marine)	Global	1
Satellite bias (land)	Global + 2.5 degree SPDE	1 + 40,962
Satellite bias (ice)	Hemispheric + 2.5 degree SPDE*	2 + 40,962

Extensions and projects in progress

- (w Liam Llamazares Elias) Penalised complexity priors for non-stationary models
- Simplified support for aggregated data models, where the predictor expression may involve integration across space (with Man Ho Suen, Andy Seaton)
- Related work (with Christopher Merchant and Xue Wang):
Multi-band satellite data with nadir and oblique views, with non-rectangular "pixels".

$$E(\text{measured}(\text{pixel}, \text{band})) = \left(\frac{1}{|D_{\text{pixel}}|} \int_{D_{\text{pixel}}} \text{conversion}[\text{SST}(\mathbf{s}), \text{TCWV}(\mathbf{s}), \text{band}]^b d\mathbf{s} \right)^{1/b}$$

- Both SST and TCWV are unknown spatial fields and b is an unknown parameter
- The "conversion" function is a deterministic function evaluated on a grid of SST and TCWV for each frequency band
- Can be implemented with numerical integration for each pixel, and spline interpolation of the conversion function

Extensions and projects in progress

- (Victor Medina) Joint covariate&outcome models for longitudinal credit risk
- (Francesco Serafini) Hawkes processes for earthquake forecasting; self-exciting Poisson processes with $\lambda(\mathbf{s}, t) = \mu(\mathbf{s}, t, \mathbf{u}) + \sum_{i; t_i < t} h(\mathbf{s} - \mathbf{s}_i, t - t_i, \mathbf{u})$ which is not log-linear.
- Copulas and transformation models; can handle non-Gaussian parameter priors as latent variables, e.g. $\lambda \sim \text{Exp}(\gamma)$ is equivalent to $\lambda = -\log[1 - \Phi(u)]/\gamma$, where $u \sim \text{N}(0, 1)$
- Extending the supported set of R-INLA models (survival models, etc)
- (w Andy Seaton) Added `sf` and `terra` support to prepare for the retirement of the `rgdal` package in 2023
- Converting the SPDE meshing code to a separate `fmesh` package
- Direct support for non-separable space-time models (INLA`spacetime`, with Elias Krainski, David Bolin, Haakon Bakka, and Haavard Rue)
- Improved support for factors and fixed effects interaction models

Further work

- How accurate are the linearised posteriors? Need diagnostic metrics for all models. Options that are more or less computable in practice include
 - $E_{\mathbf{u} \sim \tilde{p}(\mathbf{u}|\mathbf{y})}(\|\bar{\boldsymbol{\eta}} - \tilde{\boldsymbol{\eta}}\|^2)$
 - $\sum_i E_{\mathbf{u} \sim \tilde{p}(\mathbf{u}|\mathbf{y})}(|\bar{\eta}_i - \tilde{\eta}_i|^2) / \text{Var}_{\mathbf{u} \sim \tilde{p}(\mathbf{u}|\mathbf{y})}(\bar{\eta}_i)$
 - $E_{\mathbf{u} \sim \tilde{p}(\mathbf{u}|\mathbf{y})} \left(\log \left(\frac{\tilde{p}(\mathbf{u}|\mathbf{y}, \boldsymbol{\theta})}{\bar{p}(\mathbf{u}|\mathbf{y}, \boldsymbol{\theta})} \right) \right)$
- Improved convergence diagnostics and detection of unintended incorrect user input
- Interoperability with posterior analysis and plotting packages

References

- F. Lindgren, H. Rue and J. Lindström (2011),
An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach (with discussion), JRSSB, 73(4):423–498. Code available in R-INLA, see <http://r-inla.org/>
- Fabian E. Bachl, Finn Lindgren, David L. Borchers, and Janine B. Illian (2019)
inlabru: an R package for Bayesian spatial modelling from ecological survey data, Methods in Ecology and Evolution, 10(6):760–766.
<https://doi.org/10.1111/2041-210X.13168>
- CRAN package: inlabru
<https://inlabru.org/>
<https://inlabru-org.github>
<https://github.com/inlabru>
- inlabru: The Scottish INLA interface

