`spde2`: precision structure model adapted to simple SPDEs

Finn Lindgren

August 2013 (updated April 2020)

Abstract

Documentation the internals of the `spde2` SPDE model interface in R-INLA. See Lindgren et al (2011, `https://doi.org/10.1111/j.1467-9868.2011.00777.x`) for the theoretical details.

The GMRF has n nodes. The parameter vector $\boldsymbol{\theta}$ has p elements, here indexed $1, \ldots, p$.

Basis matrices

There are three basis function matrices, indexed by $k = 0, 1, 2$:

$$\boldsymbol{B}^{(k)} : n\text{-by-}(p+1) \text{ matrix}$$

The first column of each matrix is an intercept; indexing the matrix columns from 0 to p, the resulting linear combinations are

$$\phi_i^{(k)} = \boldsymbol{B}_{i,0}^{(k)} + \sum_{j=1}^{p} \boldsymbol{B}_{i,j}^{(k)} \boldsymbol{\theta}_j, \quad i = 1, \ldots, n$$

The C-interface uses the general definition, but the R-wrapper-interface has the following shortcuts:
For each of the basis matrices, when all rows are equal, only the first row needs to be supplied (all stationary SPDE models). When a basis matrix is non-zero only for the intercept, only the first row needs to be supplied (fixed model part, such as known oscillations). (I.e. for the combination of these, a single value means a constant intercept, and no parameter dependence.)

For most models, $\boldsymbol{B}^{(2)}$ will be a scalar $= 1$, so that $\phi_i^{(2)} \equiv 1$, and the link function h (see below) will be the identity transformation, eliminating this part of the model ($\boldsymbol{D}^{(2)} = \boldsymbol{I}$ in the precision expression).

Parameter transformation

Define diagonal matrices

$$
\begin{aligned}
\boldsymbol{D}^{(0)} &= \mathrm{diag}(d_i^{(0)}) = \mathrm{diag}(\exp(\phi_i^{(0)})) \\
\boldsymbol{D}^{(1)} &= \mathrm{diag}(d_i^{(1)}) = \mathrm{diag}(\exp(\phi_i^{(1)})) \\
\boldsymbol{D}^{(2)} &= \mathrm{diag}(d_i^{(2)}) = \mathrm{diag}(h(\phi_i^{(2)}))
\end{aligned}
$$

where $h(\cdot)$ is a transformation function that can be redefined. Implemented transformations:

`log` A shifted log-transformation,

$$h(x) = 2\exp(x) - 1.$$

`identity` The identity transformation

$$h(x) = x.$$

For general $\boldsymbol{B}^{(2)}$ this transformation doesn't guarantee a positive definite precision matrix, so is of limited practical use for parameter estimation, but for non-oscillating models we set $\boldsymbol{B}^{(2)} = \mathbf{1}$, and $\phi_i^{(2)} = 1$, and use the identity transformation so that $\boldsymbol{D}^{(2)} = \boldsymbol{I}$.

`logit` This replicates the now ancient `spde1` interface model, with

$$h(x) = \cos\left(\frac{\pi}{1+\exp(-x)}\right).$$

Precision matrix

Sparse matrix inputs:

$\boldsymbol{M}^{(0)}$: n-by-n sparse symmetric matrix (sometimes diagonal)

$\boldsymbol{M}^{(1)}$: n-by-n general sparse matrix

$\boldsymbol{M}^{(2)}$: n-by-n sparse symmetric matrix (sometimes all zeros)

Precision:

$$
\begin{aligned}
\boldsymbol{Q} &= \boldsymbol{D}^{(0)}\left(\boldsymbol{D}^{(1)}\boldsymbol{M}^{(0)}\boldsymbol{D}^{(1)} + \boldsymbol{D}^{(2)}\boldsymbol{D}^{(1)}\boldsymbol{M}^{(1)} + (\boldsymbol{M}^{(1)})^T\boldsymbol{D}^{(1)}\boldsymbol{D}^{(2)} + \boldsymbol{M}^{(2)}\right)\boldsymbol{D}^{(0)} \\
Q_{i,j} &= d_i^{(0)}d_j^{(0)}\left(d_i^{(1)}d_j^{(1)}M_{i,j}^{(0)} + d_i^{(2)}d_i^{(1)}M_{i,j}^{(1)} + d_j^{(2)}d_j^{(1)}M_{j,i}^{(1)} + M_{i,j}^{(2)}\right)
\end{aligned}
$$

Prior for θ

The jointly Gaussian prior for $\boldsymbol{\theta}$ is specified by an expectation vector $\boldsymbol{\mu}_\theta$ and a precison matrix $\boldsymbol{Q}_\theta$.

BLC

A linear combination of the hyperparameters is defined by a *fourth* $\boldsymbol{B}$-matrix, $\boldsymbol{B}^{(LC)}$ (called `BLC`), of size m-by-$(p+1)$ for arbitrary m, and giving the marginals for the corresponding linear transformations of the θ-vector. The SPDE R-code will construct the appropriate matrix for each given problem; sometimes it will just be a concatenation of the three other basis matrices, but usually it will be something different, and helper functions will be provided to interpret the output. This allows approximate range and field variance estimation, even for the fractional α approximations, where the internal weights for the precision matrix are different from the SPDE parameters.

Summary of inputs

$$
\begin{aligned}
n &: \text{number of GMRF nodes}, \geq 1 \\
p &: \text{number of parameters}, \geq 0 \\
m &: \text{number of parameter linear combinations}, \geq 0 \\
\boldsymbol{B}^{(0)} &: n\text{-by-}(p+1)\text{ matrix} \\
\boldsymbol{B}^{(1)} &: n\text{-by-}(p+1)\text{ matrix} \\
\boldsymbol{B}^{(2)} &: n\text{-by-}(p+1)\text{ matrix} \\
\boldsymbol{B}^{(LC)} &: m\text{-by-}(p+1)\text{ matrix} \\
\boldsymbol{M}^{(0)} &: n\text{-by-}n\text{ sparse symmetric matrix} \\
\boldsymbol{M}^{(1)} &: n\text{-by-}n\text{ general sparse matrix} \\
\boldsymbol{M}^{(2)} &: n\text{-by-}n\text{ sparse symmetric matrix} \\
\boldsymbol{\mu}_\theta &: \text{vector of length } p \\
\boldsymbol{Q}_\theta &: p\text{-by-}p\text{ sparse symmetric pos.def. matrix}
\end{aligned}
$$