ESTIMATING SEAL PUP PRODUCTION IN THE GREENLAND SEA USING BAYESIAN HIERARCHICAL MODELING (ONLINE SUPPLEMENTARY MATERIAL)

MARTIN JULLUM, THORDIS THORARINSDOTTIR, AND FABIAN E. BACHL

Abstract. This online supplement ....

1. The integrated nested Laplace approximation

The integrated nested Laplace approximation (INLA) methodology proposed by Rue et al. (2009), and implemented in the R-package INLA (www.r-inla.org), allows for computationally feasible approximate Bayesian inference for latent Gaussian models. In latent Gaussian models, n univariate observations $\boldsymbol{y} = (y_1, \dots, y_n)^\top$ are assumed to be conditionally independent given m latent Gaussian variables $\boldsymbol{z} = (z_1, \dots, z_m)^\top$ and a set of hyperparameters $\boldsymbol{\theta}$. More precisely, the INLA implementation covers models of the form

$$\begin{aligned} p(\boldsymbol{y}|\boldsymbol{z}, \boldsymbol{\theta}) &= \prod_{i=1}^{n} p(y_i|\eta_i, \boldsymbol{\theta}), \text{ with } \eta_i = \sum_{j=1}^{m} c_{ij} z_j \text{ for fixed } c_{ij}, \\ p(\boldsymbol{z}|\boldsymbol{\theta}) &\sim \mathrm{N}(\boldsymbol{\mu}(\boldsymbol{\theta}), Q(\boldsymbol{\theta})^{-1}), \\ \boldsymbol{\theta} &\sim p(\boldsymbol{\theta}), \end{aligned} \tag{1.1}$$

where the latent variables $\boldsymbol{z}$ may depend on additional (fixed) covariates, for a large class of models for $\boldsymbol{y}$.

For computationally fast inference it is essential that the precision matrix $Q(\boldsymbol{\theta})$ is sparse and that the parameter vector $\boldsymbol{\theta}$ is of a fairly low dimension. This covers models where the latent field is a *Gaussian Markov random field* (GMRF). For the inference, INLA utilizes several nested Laplace approximations. That is, the posterior distribution of $\boldsymbol{\theta}$ is approximated by

$$p(\boldsymbol{\theta}|\boldsymbol{y}) \approx \tilde{p}(\boldsymbol{\theta}|\boldsymbol{y}) \propto \left. \frac{p(\boldsymbol{y}, \boldsymbol{z}, \boldsymbol{\theta})}{p_G(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})} \right|_{\boldsymbol{z}=\boldsymbol{z}^*(\theta)}, \tag{1.2}$$

where $p_G(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})$ is a Gaussian approximation to the full conditional distribution of $\boldsymbol{z}$, and $\boldsymbol{z}^*(\boldsymbol{\theta})$ is the mode of $p(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})$ for a given $\boldsymbol{\theta}$. The marginals of this low-dimensional posterior distribution are typically computed by direct numerical integration. The marginals for the latent field, $p(z_j|\boldsymbol{y})$, are typically computed by first obtaining a Laplace approximation $\tilde{p}(z_j|\boldsymbol{\theta}, \boldsymbol{y})$ similar to (1.2), or a Taylor approximation of that distribution, and then solve $\int \tilde{p}(z_j|\boldsymbol{\theta}, \boldsymbol{y})\tilde{p}(\boldsymbol{\theta}|\boldsymbol{y})\, \mathrm{d}\boldsymbol{\theta}$ by numerical integration. See Rue et al. (2009) and Martins et al. (2013) for further details.

2. Triangular mesh used in SPDE-INLA

The triangular mesh we use to fit the LGCP model with SPDE-INLA is displayed in Figure 1. To overcome boundary effects, we extend the area that is modeled quite a bit beyond the whelping region as recommended by Lindgren et al. (2011). In order to properly represent observations being aggregated over a certain spatial domain (the photos), we construct the mesh in a specific way such that there is a mesh node in the center of every photo and that the corresponding Voronoi tessellation (Watson, 1981) matches the photo as close as possible. The Voronoi tessellation is used to specify the weight or "offset" of the observations used in the Poisson distribution. Voronoi tessellations that match the photos are obtained by placing mesh nodes at the center point of each photo, and at a distance equal to the height of the photo both above and below the center point, see Figure 2. For respectively, the leftmost and rightmost photo on each transect, additional points are placed to the left and the right, at a distance equal to the width of each photo. Note that this procedure is merely a technical task carried out in order to fit the problem with the SPDE approach while still using the exact locations of the irregular lattice (the photos) in a Poisson regression formulation. The mesh has a coarser resolution elsewhere, where fine resolution detail of the latent field cannot be easily estimated.

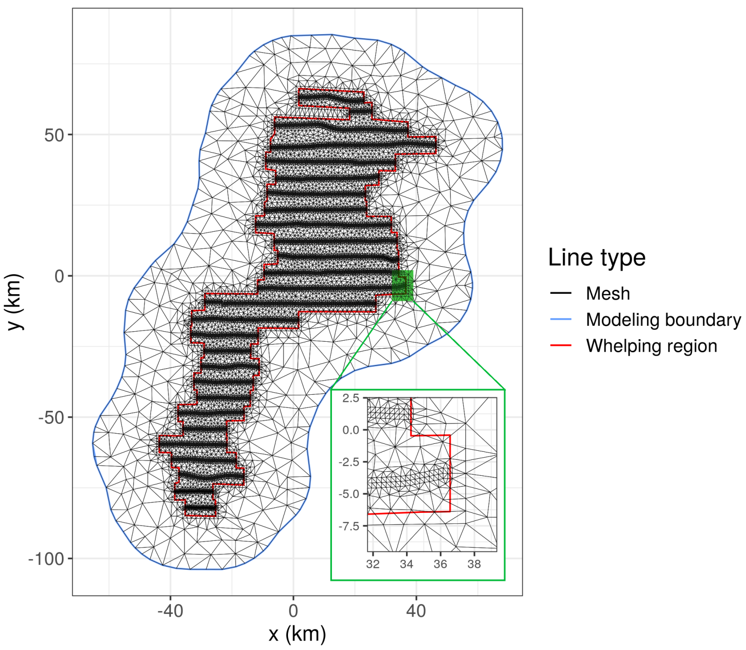

FIGURE 1. The triangular mesh used in the SPDE-INLA analysis for the LGCP model. The bottom right corner shows a zoomed in version of the mesh for the green area above.

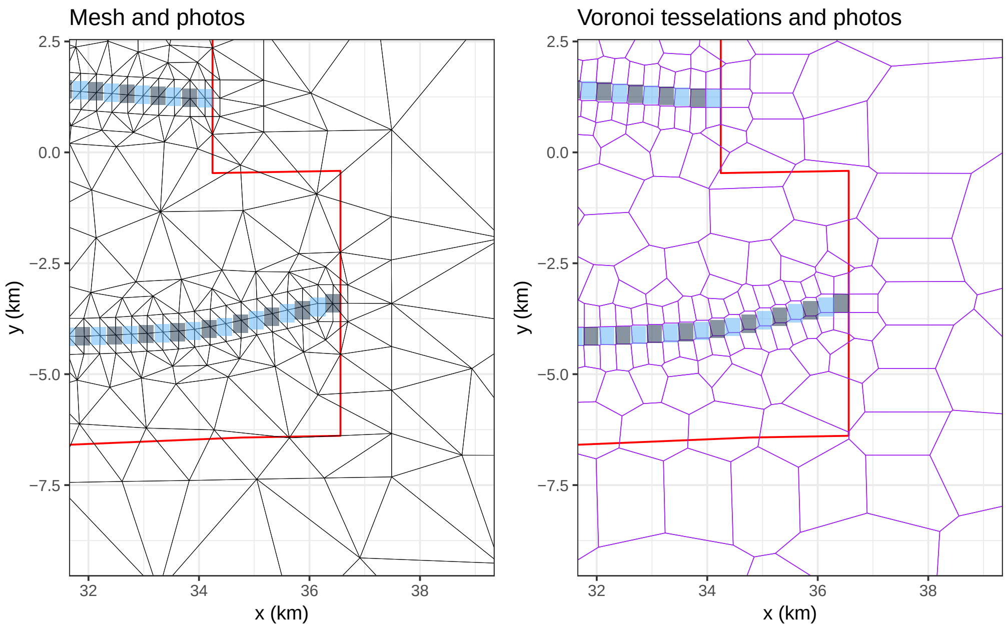

FIGURE 2. Left: The triangular mesh used in the SPDE-INLA analysis for the LGCP model (in black) with the photos (gray and blue squares) for the zoomed-in area in Figure 1. Right: The resulting Voronoi tesselations (in purple) with the photos for the same area.

REFERENCES

[1] Lindgren, F., Rue, H. and Lindström, J. (2011) An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73**, 423–498.

[2] Martins, T. G., Simpson, D., Lindgren, F. and Rue, H. (2013) Bayesian computing with INLA: New features. *Computational Statistics & Data Analysis*, **67**, 68–83.

[3] Rue, H., Martino, S. and Chopin, N. (2009) Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*, **71(2)**, 319–392.

[4] Watson, D. F. (1981) Computing the n-dimensional delaunay tessellation with application to voronoi polytopes. *The computer journal*, **24**, 167–172.

Norwegian Computing Center, Oslo, Norway.
Email address: jullum@nr.no

Norwegian Computing Center, Oslo, Norway.
Email address: thordis@nr.no

School of Mathematics, University of Edinburgh, Edinburgh, Scotland.
Email address: bachlfab@gmail.com