Norsk Regnesentral
NORWEGIAN COMPUTING CENTER

Note

Estimating seal pup production in the Greenland Sea using Bayesian hierarchical modeling

Note no SAMBA/04/18

Authors
Martin Jullum
Thordis L. Thorarinsdottir
Fabian Bachl

Date 29th April 2018

Norsk Regnesentral
NORWEGIAN COMPUTING CENTER

The authors

Martin Jullum is Research Scientist and Thordis L. Thorarinsdottir is Chief Research Scientist at the Norwegian Computing Center in Oslo, Norway. Fabian Bachl is an Analyst with d-fine in Frankfurt.

Norwegian Computing Center

Norsk Regnesentral (Norwegian Computing Center, NR) is a private, independent, non-profit foundation established in 1952. NR carries out contract research and development projects in information and communication technology and applied statistical-mathematical modelling. The clients include a broad range of industrial, commercial and public service organisations in the national as well as the international market. Our scientific and technical capabilities are further developed in co-operation with The Research Council of Norway and key customers. The results of our projects may take the form of reports, software, prototypes, and short courses. A proof of the confidence and appreciation our clients have in us is given by the fact that most of our new contracts are signed with previous customers.

Norsk Regnesentral
Norwegian Computing Center
Postboks 114, Blindern
NO-0314 Oslo, Norway

Besøksadresse
Office address
Gaustadalléen 23
NO-0373 Oslo, Norway

Telefon · telephone
(+47) 22 85 25 00
Telefaks · telefax
(+47) 22 69 76 60

Internett · internet
www.nr.no
E-post · e-mail
nr@nr.no

Title	**Estimating seal pup production in the Greenland Sea using Bayesian hierarchical modeling**
Authors	**Martin Jullum** <jullum@nr.no> **Thordis L. Thorarinsdottir** <thordis@nr.no> **Fabian Bachl** <bachlfab@googlemail.com>
Date	29th April 2018
Publication number	SAMBA/04/18

Abstract

The Greenland Sea is an important breeding ground for harp seals (Pagophilus groenlandicus) and hooded seals (Cystophora cristata). An estimate of the annual seal pup production is a critical factor in the abundance estimation needed for management of the species. Estimates of seal pup production are usually based on counts from aerial photographic surveys. However, due to the large extent of typical whelping regions, only a minor part of the complete area can be photographed. To estimate the total seal pup production, we propose a Bayesian hierarchical modelling approach motivated by viewing the seal pup appearances as a realization of a log-Gaussian Cox process using covariate information from satellite imagery as a proxy for ice-thickness. For inference, we utilize the spatial partial differential equation (SPDE) module of the integrated nested Laplace approximation (INLA) framework. In a case study using survey data from 2012, we compare our results with existing methodology in a comprehensive cross-validation study. The new proposed method improves local estimation performance and more accurately addresses the associated uncertainty.

Keywords	Spatial point process; Probabilistic prediction; Uncertainty quantification
Target group	Statisticians, ecologists
Availability	Open
Project	PointProcess
Project number	220708
Research field	Statistics
Number of pages	24

1 Introduction

Three stocks of harp seals (Pagophilus groenlandicus) and two (possibly three) stocks of hooded seals (Cystophora cristata) inhabit the North Atlantic Ocean where they have been harvested for centuries (Kovacs and Lavigne, 1986; Sergeant, 1974, 1991). Monitoring the abundance of seals is vital for controlling the biodiversity in the region. State-of-the-art seal population models are dynamically built based on historical catch data (Øigård et al., 2014a,b). The main ingredient in these models is the total pup production in a given year which needs to be quantified based on on-site observational data since other quantification methods based on catch-at-age and mark-recapture data etc. are considered unreliable (ICES, 2014). The whelping regions in the North Atlantic typically cover several thousand square kilometers so that the total pup production needs to be estimated based on observations from a minor part of the region. The estimated total pup production and the associated uncertainty then enter the dynamic population model (Øigård et al., 2010).

The seal pup count data are typically obtained by manual counting based on photographs stemming from an aerial photographic survey conducted by flying along transects sparsely covering the whelping region. The survey methodology is discussed in more detail in Section 2. The traditional method for estimating the total pup production based on such count data is that of Kingsley et al. (1985) which assumes a homogeneous dispersion of seals across the entire whelping region. Salberg et al. (2009) propose a generalized additive modeling (GAM) approach (Hastie and Tibshirani, 1990), assuming the counts follow a negative binomial distribution and taking the spatial location of the counts into account. For data that is close to homogneous, the negative binomial GAM approach and the Kingley method yield similar estimates. However, the Kingley method may possess a positive bias when the spatial distribution of the pups is clustered (Øigård et al., 2010; Salberg et al., 2008). Additionally, the GAM method produces much smaller uncertainty bounds than the homogenous Kingley approach.

In this paper, we propose a new method for estimating the total seal pup production. We view the seal pup appearances as a spatial point process (Møller and Waagepetersen, 2003) and model the point pattern of the seal pups as a log-Gaussian Cox process (LGCP; Møller et al., 1998). with a spatial latent field which also allows additional covariate information to be accounted for. In a Bayesian formulation with priors on the model parameters, the seal pup abundance estimate is represented by the posterior predictive distribution found by integrating the posterior distribution over the spatial domain of the whelping region, instead of a single point estimate accompanied with a variance estimate. This Bayesian hierarchical model can be fitted by utilizing the Spatial Partial Differential Equation (SPDE) approach of the Integrated Nested Laplace Approximation (INLA) (Lindgren et al., 2011; Rue et al., 2009). The final posterior predictive distribution can subsequently be computed from this fitted model by a sampling approach.

To illustrate and test this methodology, we use seal pup photo counts from an aerial

photographic survey in the Greenland Sea in March 2012 with two different types of seals, harp and hooded seals. The data set contains the spatial location of each photo and the corresponding pup count. To be more informative about the non-observed areas, we include covariate information extracted from satellite imagery captured on the very same date as the aerial photographic survey was conducted, to act as a proxy for ice thickness. Compared to the two other procedures our method gives larger uncertainties, especially for the harp seals. To validate these differences, we compare our proposed method with the two reference methods in a comprehensive validation scheme. The scheme suggests our method performs best on local level, and comparable on a more global scale. Further investigations suggest that the larger uncertainty in our method is indeed more realistic.

The rest of the paper is organized as follows: Section 2 describes the surveying method used to gather seal pup observational data, in addition to specific details related to our particular seal pup data set. The satellite imagery and the covariate information extracted therefrom, are also discussed. Section 3 is the main section of the paper. It starts by giving some background to our modeling approach through brief introductions to the relevant parts of point processes and the INLA framework, before going into detail about the present modeling approach. Subsequently, we discuss the two reference approaches in more detail, before we describe the validation scheme we use to verify our procedure and compare it to the reference approaches. Section 4 contains the results from our study, including both specifics of the model fitted with our procedure, and the validation results. Section 5 contains some concluding remarks and pointers to further work.

2 Data

In this section we describe the surveying method and additional modeling information we have obtained through the satellite imagery.

2.1 Surveying method

Before conducting the aerial photographic survey with the purpose of monitoring the seal pup abundance, the marine researchers typically perform a helicopter reconnaissance survey. An important task when performing these reconnaissance surveys is to locate the patches where the seals live. This is important in order to limit the survey area for the more expensive airborne photographic survey. The actual photographic survey is conducted by flying a survey aircraft equipped with advanced photographic equipment and GPS along a number of fixed distanced transects sparsely covering the survey area.

In this particular survey in March 2012, the airplane flew at an altitude of about 330m, and took totally 2792 photos along 27 transects, each covering 226×346m of ground level. Due to fog, an exception was made for the two southernmost transects, which were flown at altitude 250m, with photos covering 170×260m. In any case, along each transect, the cameras were turned on when the first seal was spotted from the airplane, and kept on taking photos continuously until the ice edge (which no seals can live outside) was

reached on the eastern side, and when no seals were spotted for an extended period to the west. A consequence of this survey setup is that even if no photos were taken to the west and east of the transect lines, the survey method implies that there are no seals there. The whelping region is approximately defined as the union of the 1.5Nm (≈ 2.8km) bands around each transect. Thus, when computing the posterior predictive distributions with the fitted models, we do only count predictions within the whelping region. More details about this survey may be found in Øigård et al. (2014a,b).

2.2 Seal pup counts

Following the airborne survey described above, experienced marine researchers carefully examine each captured photo and note how many seal pups of each species there are in each photo. Quality checks with multiple examinations also limit the measurement error introduced in this step (Øigård et al., 2014a,b). Our seal count data set contains the coordinates and extent of each photo, in addition to the number of seal pups of each species observed in them. These data are plotted in Figure 1 along with the flown transects and extent of the whelping region. As seen from the figure, there tends to be more seal types clustered towards the middle eastern boundary and southern corner of the whelping region.

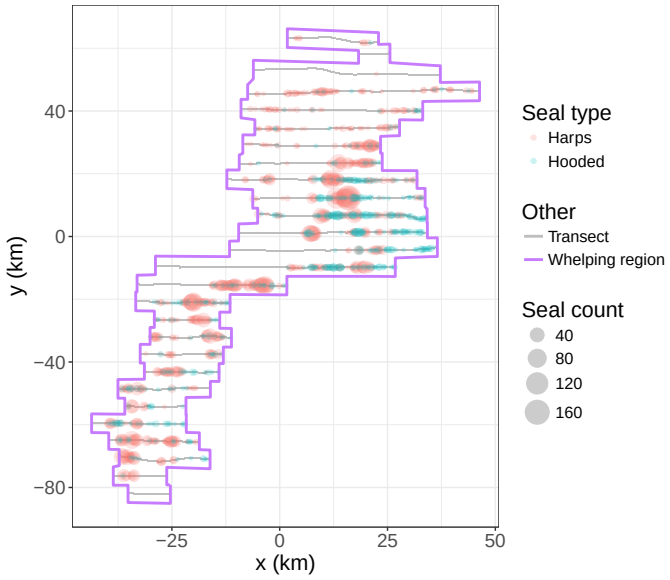

Figure 1. Harp and hooded seal pup data from the Greenland 2012 survey.

2.3 Satellite imagery

As seals cannot live where there are no large ice floes to lie on, information about which areas are covered by ice floes and which areas are merely open water, is highly relevant when estimating the seal pup abundance. In an attempt to account for this, we have collected high resolution satellite imagery from the whelping region captured the very same day as the airborne photographic survey was conducted. From this satellite imagery we have extracted a variable which acts as a proxy for the ice thickness. This density variable is displayed in Figure 2. Comparing the satellite data to the seal pup counts in Figure 1,

we see that the seal pup counts appears to be higher in the areas with high ice density, than in the areas with lower ice density.

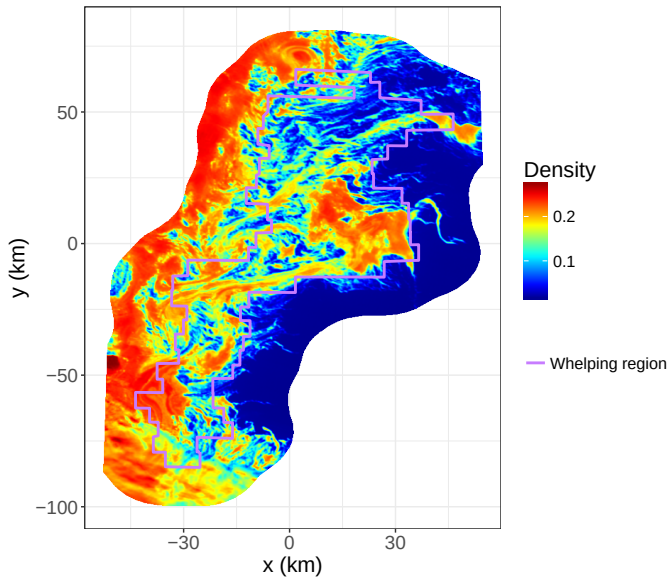

Figure 2. Satellite data used as covariates in the model fitting process.

3 Methods

This section starts with a brief introduction to point processes, and the key components of INLA and the SPDE approach which makes our proposed method computational feasible. We then describe the details of our proposed modeling approach, the two references methods we compare our method to, and finally the verification scheme used to validate and compare the methods.

3.1 Background

3.1.1 Point processes

A spatial stochastic point process is a mathematical description of the random process at which points or locations are distributed in space. The collection of such observations is called a point pattern. In informal mathematical terms, which typically are sufficient for practical purposes, a spatial point process Y is a random collection of points in a bounded observation region $A \subset \mathbb{R}^2$. By random collection we mean that both the number of points and their actual location are random. The most fundamental type of point process is the *Poisson point process*. This process is specified by a single intensity function $\lambda : A \mapsto [0, \infty)$ which determines the density of points in any location in A, and has the property that the number of points, say $N(B)$, of any Borel set B is Poisson distributed with mean $\mu(B) = \int_B \lambda(\boldsymbol{s})\, \mathrm{d}\boldsymbol{s}$, i.e. $N(B) \sim \text{Pois}(\mu(B))$, and independent of $N(B^*)$ for any other non-overlapping Borel set B^*.

The Cox process introduced by Cox (1955) is a generalization of the Poisson point pro-

cess where the intensity function λ is itself stochastic. A popular special case of this hierarchical model is the Log-Gaussian Cox process (LGCP) (Møller et al., 1998), where the intensity is assumed to be log-Gaussian; i.e. there is a Gaussian latent field Z, and given $\lambda = \exp(Z)$, Y is a Poisson point process with intensity λ. This type of model is indeed very flexible and ought to fit well with a great variety of natural processes.

The very nature of point processes, having a model dimensionality which changes with the number of data points, make them difficult to work with both in theory and in practice. Advances within computational statistics the last decades have however made complicated spatial point process models more accessible to applied researchers, see e.g. Møller and Waagepetersen (2003).

3.1.2 Integrated nested Laplace approximation

The integrated nested Laplace approximation (INLA) methodology by Rue et al. (2009), and the accompanying R-package `INLA` allows for computationally feasible approximate Bayesian inference for a class of models called latent Gaussian models. That is, hierarchical models where n univariate observations gathered in $\boldsymbol{y} = (y_1, \ldots, y_n)^\top$, is assumed to be conditionally independent given m latent Gaussian variables gathered in $\boldsymbol{z} = (z_1, \ldots, z_m)^\top$, and a set of hyperparameters $\boldsymbol{\theta}$. More precisely, the `INLA` implementation covers models of the form

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

where the latent variables in $\boldsymbol{z}$ may also depend on additional (fixed) covariates when appropriate. For computationally fast inference it is essentially that $Q(\boldsymbol{\theta})$ is sparse and that the model parameter vector $\boldsymbol{\theta}$ is of a fairly low dimension. This covers models where the latent field is a Gaussian Markov random field (GMRF), which is often used within spatial statistics. The basic principle of the INLA methodology is, as the name reveals, to utilize several nested Laplace approximations. The posterior distribution of $\boldsymbol{\theta}$ is approximated using the Laplace approximation

$$
p(\boldsymbol{\theta}|\boldsymbol{y}) \approx \tilde{p}(\boldsymbol{\theta}|\boldsymbol{y}) \propto \left. \frac{p(\boldsymbol{y}, \boldsymbol{z}, \boldsymbol{\theta})}{p_G(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})} \right|_{\boldsymbol{z}=\boldsymbol{z}^*(\theta)}, \tag{1}
$$

where $p_G(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})$ is a Gaussian approximation to the full likelihood 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, $p(x_j|\boldsymbol{y})$, are typically computed by first obtaining $\tilde{p}(z_j|\boldsymbol{\theta}, \boldsymbol{y})$ with a Laplace approximation of $p(\boldsymbol{z}_{-j}|z_j, \boldsymbol{\theta}, \boldsymbol{y})$ similar to (1), or a Taylor approximation of that distribution; and then to use the resulting approximation, along with $\tilde{p}(\boldsymbol{\theta}|\boldsymbol{y})$ to numerically integrate out $\boldsymbol{\theta}$. See Martins et al. (2013); Rue et al. (2009) for details.

The spatial partial differential equation (SPDE) approach of Lindgren et al. (2011) has made it possible to perform computationally efficient Bayesian inference with INLA also

for a class of *continuous* latent fields. This is made possible by (approximately) transforming the continuous latent field to a certain GMRF, formulated through the solution of a SPDE, when doing the computations. The key ingredient in this approach is to approximate the continuous field $Z_0(\boldsymbol{s})$ by a field $Z(\boldsymbol{s})$ living on a triangular mesh. For a triangular mesh with m triangle vertices, we write

$$Z(\boldsymbol{s}) = \sum_{j=1}^{m} z_j \phi_j(\boldsymbol{s}), \tag{2}$$

where $\boldsymbol{z} = (z_1, \ldots, z_m)^\top$ is a multivariate Gaussian random vector and $\{\phi_j(\boldsymbol{s})\}_{j=1}^m$ is a set of deterministic linearly independent basis functions which are piecewise linear between the vertices and chosen such that $\phi_j(\boldsymbol{s})$ is 1 at vertex j, and 0 at all other vertices. A consequence of the representation in (2) is that $Z(\boldsymbol{s})$ is fully determined by $\boldsymbol{z}$: $Z(\boldsymbol{s})$ takes the value z_j at vertex j while its values inside the triangles are determined by linear interpolation. Let us further equip $Z(\boldsymbol{s})$ with the Matèrn covariance function

$$\text{Cov}(Z(\boldsymbol{s}), Z(\boldsymbol{t})) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} (\kappa \|t - s\|)^\nu K_\nu(\kappa \|t - s\|), \tag{3}$$

where $\nu > 0$ is a smoothing parameter, K_ν is the modified Bessel function of the second kind, $\kappa > 0$ is a scaling parameter and σ^2 is the marginal variance. Lending on this type of latent field being a solution to a certain SPDE, the precision matrix Q of $\boldsymbol{z}$ takes an analytical form which can be approximated by a sparse matrix $\tilde{Q}$. Since $Z(\boldsymbol{s})$ is completely determined by the $\boldsymbol{z}$, this allows continuous field computations to be carried out approximately using this GMRF as for the discretely indexed latent Gaussian models, using e.g. the INLA software. Note that following certain guidelines for constructing the triangular mesh, the approximation error involved in using this GMRF representation, is typically small (Lindgren et al., 2011; Simpson et al., 2012).

For a complete introduction and review of the INLA framework, including the SPDE approach, see the reference book Blangiardo and Cameletti (2015) and the review paper Rue et al. (2016).

3.2 The present modeling approach

Before any observations are made, both the number of seal pups and where they are located within the whelping region, are truly random. Despite this being the precise nature of a point process, such an approach has not, to the best of our knowledge, been used to study abundance or other aspects of seal wildlife. With knowledge of the locations of the seal mothers, and viewing the number of pups produced by each seal mother as a Poisson distributed variable, the Poisson point process is the natural modeling approach. However, as the locations of the seal mothers is unknown to us, it is natural to consider the intensity as a random component as well, i.e. introducing a latent field, brings us to the very flexible Cox process (also known as the doubly stochastic Poisson point process). Assuming a log-Gaussian distribution for the latent field gives the frequently applied Log-Gaussian Cox process and enables us to take advantage of the computationally efficient INLA framework.

A complicating component of the present estimation problem is that the exact positions of the counted seal pups are not available. Instead, the seal pup counts are provided as aggregated counts per photo. Thus, our data are merely point process data aggregated to an irregular lattice, as opposed to actual point patterns. This disallows us from directly applying the approach of Simpson et al. (2016), which fits LGCPs using the SPDE-INLA framework by directly approximating the likelihood function of the LGCP. A method which is accessible for both actual point patterns and for aggregated data, is that of Rue et al. (2009, Section 5.5), which suggests approximating the LGCP by fitting a Poisson regression formulation to aggregated counts on a regular lattice constituting a GMRF. However, as the photos do not compose a regular lattice, such an approach would call for further, typically quite rough, approximations.

The modeling approach we have chosen for this situation is a Poisson regression formulation based on aggregations from a log-Gaussian Cox process, exactly as in Rue et al. (2009, Section 5.5). However, instead of requiring the lattice to be regular, we work with a triangulated irregular lattice allowing us to apply the SPDE procedures of INLA. This is preferable over the regular lattice approach of Rue et al. (2009, Section 5.5) for two reasons. First, we may work directly with the data exactly as provided to us. Secondly, contrary to the regular lattice approach, the latent field is continuous which gives a more precise representation of the phenomenon.

The continuous domain latent field we use in the modeling takes the form

$$Z(\boldsymbol{s}) = \alpha + \boldsymbol{\beta}^\top \boldsymbol{x}(\boldsymbol{s}) + f(\boldsymbol{s}), \tag{4}$$

where α is an intercept term, $\boldsymbol{\beta}$ is a coefficient vector corresponding to $\boldsymbol{x}(\boldsymbol{s}) = (q(\boldsymbol{s}), s_1, s_2, s_1 s_2)^\top$, where $q(\boldsymbol{s})$ contains the density variable from the satellite imagery, while the rest are components modeling linear spatial effects. Finally $f(\boldsymbol{s})$ is a SPDE based continuous Gaussian field with the Matèrn covariance function given in (3) taking care of the non-linear spatial dependence. To make the intercept identifiable, we restrict $f(\boldsymbol{s})$ to integrate to zero over the modeling region. Due to poor identifiability, we fix the Matèrn smoothing parameter $\nu = 2$. This is also the default in the INLA software. The other hyperparameters related to the latent field are equipped with default priors specified by the INLA software. For our mesh these are $\theta_1 \sim \mathrm{N}(1.328, 10)$, $\theta_2 \sim \mathrm{N}(-2.594, 10)$ and independent, where $\theta_1 = \log(\tau), \theta_2 = \log(\kappa)$ and $\sigma^2 = 1/(4\pi\kappa^2\tau^2)$. In addition α is assigned the improper prior $\alpha \sim N(0, \infty)$, while $\boldsymbol{\beta} \sim N(0, 1000\mathcal{I}_4)$, where $\mathcal{I}_4$ denotes the 4×4 dimensional identity matrix. As all these priors are vague, they will in very limited degree influence the final results.

The triangular mesh used in the modeling is displayed in Figure 3. Notice that the mesh is detailed around the photos and wider elsewhere, where fine detail of the latent field is anyway impossible to obtain. To overcome boundary effects, we extend the area that is modeled quite a bit beyond the whelping region. This is common strategy working with these, see Lindgren et al. (2011).

To pan out how we do the actual fitting of our data, and produce a posterior predictive distribution for the seal pup abundances, let $N(A_i), i = 1, \ldots, n$ denote the number of

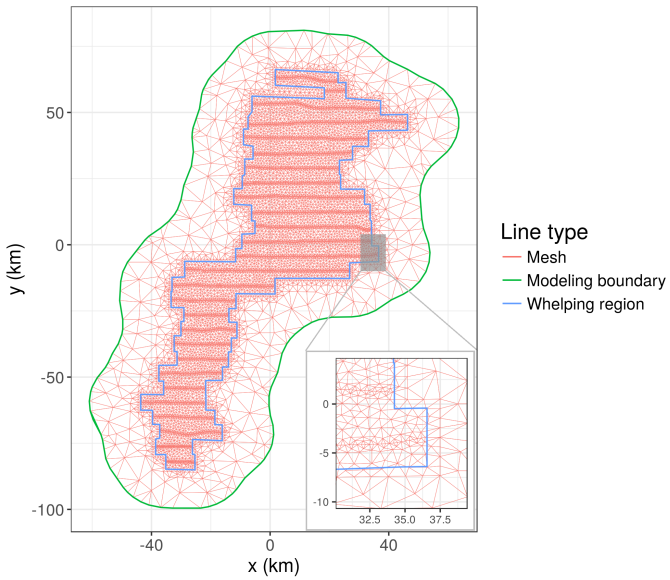

Figure 3. Mesh used in the modeling. The bottom right corner shows a zoomed in version of the mesh for the grey area above.

seal pups in each of the $n = 2792$ photos with domains $A_i, i = \ldots, n$. As per the LGCP formulation, conditional on the latent field Z, the number of seal pups in photo i, $N(A_i)$, is Poisson distributed with parameter $\lambda = \int_{A_i} \exp(Z(\boldsymbol{s}))\,\mathrm{d}\boldsymbol{s}$, i.e.

$$N(A_i)|Z \sim \text{Pois}(\lambda = \int_{A_i} \exp(Z(\boldsymbol{s}))\,\mathrm{d}\boldsymbol{s}), i = 1, \ldots, n \tag{5}$$

As this is precisely the format of our data, we use the Poisson likelihood in (5) to fit our model using the `INLA` R-package. Let now Ω denote the whelping region, i.e. the area where we would like to estimate the abundance, $N(\Omega)$. We shall use the posterior predictive distribution within Ω as the full Bayesian assessment of the seal pup abundance. This quantity is defined as

$$\begin{aligned} p(N(\Omega)|N(A_1), \ldots, N(A_n)) &= \int p(N(\Omega), Z|N(A_1), \ldots, N(A_n))\,\mathrm{d}Z \\ &= \int p(N(\Omega)|Z)p(Z|N(A_1), \ldots, N(A_n))\,\mathrm{d}Z. \end{aligned} \tag{6}$$

Having fitted the model in (5), with the latent field in (4), its aforementioned priors, and the mesh in (3), the INLA software produces posterior distributions for all individual hyperparameters, and enables sampling from the posterior of the complete latent field $p(Z|N(A_1), \ldots, N(A_n))$. This allows us to use a Monte Carlo approximation to the integral in (6):

$$\int p(N(\Omega)|Z)p(Z|N(A_1), \ldots, N(A_n))\,\mathrm{d}Z \approx \frac{1}{K}\sum_{k=1}^{K} p(N(\Omega)|\tilde{Z}_k),$$

where $\tilde{Z}_k$ is the k-th sample of the posterior latent field. Further, by the point process definition, $p(N(\Omega)|Z = \tilde{Z}_k) \sim \text{Pois}(\lambda = \int_\Omega \tilde{Z}_k(\boldsymbol{s})\,\mathrm{d}\boldsymbol{s})$. The integral here can be solved by

e.g. a simple Riemann midpoint rule by dividing Ω into J rectangles $B_1, \ldots, B_J$ centered in $s_1, \ldots, s_J$, and using $\int_\Omega \tilde{Z}_k(\boldsymbol{s})\,\mathrm{d}\boldsymbol{s} \approx \sum_{j=1}^{J} \tilde{Z}_k(s_j)|B_j|$. The final approximation to the posterior predictive distribution is

$$p(N(\Omega)|N(A_1), \ldots, N(A_n)) \approx \frac{1}{K}\sum_{k=1}^{K} \text{Pois}(\lambda = \sum_{j=1}^{J} \tilde{Z}_k(s_j)|B_j|), \tag{7}$$

i.e. a Poisson mixture distribution.

3.3 Reference approaches

This subsection gives a brief outline of existing methodology for estimating the seal pup abundance.

3.3.1 Kingsley's method and the homogenous Poisson model

The conventional method used to estimate seal pup abundance based on aerial photographic transect surveys like those we have here, is what we shall refer to as Kingsley's method. This method, due to Kingsley et al. (1985), is fundamentally simple: For each transect $T_1, \ldots, T_{27}$ covering the space $A_{T_k} = \bigcup_{\{A_i \text{ belongs to transect } T_k\}} A_i$, compute the seal pup count $N_{T_k} = \sum_{A_j \in T_k} N(A_j)$. Then the estimate of the total number of seal pups is

$$\widehat{\mu}_{\text{Kingsley}}(N(\Omega)) = \frac{|A_\Omega|}{\sum_{k=1}^{27}|A_{T_k}|}\sum_{k=1}^{27} N_{T_k}. \tag{8}$$

Kingsley et al. (1985) also provide an estimate of the variance related the abundance estimate, based on serial differences between the transects, which Salberg et al. (2008) provides a modification to.

Since this non model based method does not provide predictive distribution, but only a point estimate and variance estimate, it is hard to properly compare this method with our suggested Bayesian procedure. We will therefore not perform validation tests, as described in Section 3.4, against this method.

Kingsley's method is built on the principle of homogeneity within the whelping region. As such, a Poisson likelihood model with a homogenous intensity λ_0 is therefore quite similar to this approach. Using the photo area $|A_i|$ as offset, one then models $N(A_i) \sim \text{Pois}(\lambda = |A_i|\lambda_0)$, which can be fitted with standard software for generalized linear models. With this model, then the posterior predictive distribution for the seal pup abundance becomes

$$p_{\text{Hom.Pois}}(N(\Omega)|N(A_1), \ldots, N(A_n)) = \text{Pois}(\lambda = |\Omega|\widehat{\lambda}_0), \tag{9}$$

where $\widehat{\lambda}_0$ is the Maximum Likelihood (ML) estimate of λ_0 based on the photos. (9) does not take the uncertainty involved in estimating λ_0. To account for this, we rely on the asymptotic distribution of the ML estimator

$$N(\lambda_0, \text{Var}(\widehat{\lambda}_0)), \tag{10}$$

where $\mathrm{Var}(\widehat{\lambda}_0)$ is the asymptotic variance of $\widehat{\lambda}_0$, generally available as a closed from expression. Similar in spirit to (7), this yields

$$p_{\text{Hom.Pois}}(N(\Omega)|N(A_1),\ldots,N(A_n)) \approx \frac{1}{K}\sum_{k=1}^{K}\text{Pois}(\lambda = |\Omega|\tilde{\lambda}_k), \tag{11}$$

where $\tilde{\lambda}_k$ is k-th sample from (10). In the verification process, we will use the predictive distribution in (11) in place of Kingsley's method.

3.3.2 Generalized Additive Model

The Kingsley and related methods described above are homogenous in the sense that they implicitly assume that the seal pup intensity is constant in the whelping region. Such an assumption is typically not valid. Salberg et al. (2009) tries to overcome this weakness by modeling the seal pup abundance by a generalized additive model based on a negative binomial likelihood. Denoting by $\text{NegBin}(\mu,\kappa)$ the parametrization of the negative binomial distribution with mean μ and shape κ, this method takes the form

$$N(A_j) \sim \text{NegBin}(\mu = |A_j|\exp(S(\boldsymbol{s}_j)),\kappa), \tag{12}$$

where $|A_j|$ as before is the non-estimated offset, and with $S(\cdot)$ is a spatial smoothing component. The spatial smoothing component $S(\cdot)$ is a so-called thin-plate smoothing regression spline which fits overlapping cubic regression on a set of artificial knots in space, producing a smooth, nonlinear spatial effect. A generalized cross validation (GCV) criterion is used to select the right amount of smoothing. See Øigård et al. (2010); Salberg et al. (2009) for details about the model. To compare this modeling procedure with our suggestion under fair circumstances, we extend the GAM approach with a covariate effect β_q for the satellite image density variable $q(\boldsymbol{s})$. That is, we extend (12) to

$$N(A_j) \sim \text{NegBin}(\mu = |A_j|\exp(S(\boldsymbol{s}_j) + \beta q(\boldsymbol{s}_j)),\kappa).$$

There is a wide range of software available for fitting such GAMs, we use the `gam` function in the `mgcv` R-package. In contrast to the Poisson distribution, the negative binomial distribution is not closed under additivity. That is, the sum of two negative binomial distributions with different means does not possess a negative binomial distribution. Thus, to arrive at a full posterior predictive distribution for the seal pup abundance, we need to rely on an underlying conditional independence condition stating that conditional on $S(\cdot),\beta$ and κ, then $N(A_j)$ and $N(A_k), k \neq j$ are independent. Using this property, the posterior predictive distribution for the seal pup abundance with the GAM model can be written

$$p_{\text{GAM}}(N(\Omega)|N(A_1),\ldots,N(A_n)) = \frac{1}{J}\sum_{j=1}^{J}\text{NegBin}(\mu = |B_j|\exp(\widehat{S}(\boldsymbol{s}_j) + \widehat{\beta}q(\boldsymbol{s}_j)),\widehat{\kappa}), \tag{13}$$

where $B_j, j = 1,\ldots,J$ is a gridding of Ω of size similar to the photos, centered in $\boldsymbol{s}_j, j = 1,\ldots,J$. To account for the uncertainty in estimating $\widehat{S}(\cdot)$ and $\widehat{\kappa}$, we use the estimators'

asymptotic normal distribution, as in (10) and (11) for the homogenous Poisson model, to arrive at

$$p_{\text{GAM}}(N(\Omega)|N(A_1),\ldots,N(A_n)) \approx \frac{1}{JK}\sum_{k=1}^{K}\sum_{j=1}^{J}\text{NegBin}(\mu = |B_j|\exp(\tilde{S}_k(\boldsymbol{s}_j) + \tilde{\beta}_k q(\boldsymbol{s}_j)), \tilde{\kappa}_k), \tag{14}$$

where $\tilde{S}(\cdot)_k, \tilde{\beta}_k$ and $\tilde{\kappa}_k$ are k-th samples of respectively the smooth spatial effect, the covariate effect, and the scaling parameter.

Salberg et al. (2009) relies on the negative binomial distribution instead of the Poisson distribution for the GAM modeling by arguing that the latter is inappropriate due to overdispersion. We believe this argument is only valid with a homogenous intensity, and therefore add the same model, but with a Poisson distribution, to our list of models to compare. That is, we also fit the following GAM Poisson model

$$N(A_j) \sim \text{Poisson}(\lambda = |A_j|\exp(S(\boldsymbol{s}_j) + \beta q(\boldsymbol{s}_j))). \tag{15}$$

3.4 Verification

To validate our suggested procedure and compare its performance to the reference methods outlined in Section 3.3.1 and 3.3.2, we go through a fairly comprehensive (cross) validation scheme. We rely on two procedures for subsetting the data. The former method is a standard 10-fold cross validation setup, where we randomly remove 10% of the photos each time, such that each photo is removed exactly once. In the latter we remove all photos in one full transect at a time, such that each transect is removed exactly once, leaving us with 27 different subsets. For both methods we fit the complete model for every subset and compute posterior predictive distributions for every photo that is removed along with posterior predictive distributions for the sum of the removed transects (corresponding to the full transect for the latter method). To measure the quality of the posterior predictive distributions, we use two performance measures: the logarithmic score and the Continuous Rank Probability Score (CRPS) (Matheson and Winkler, 1976). Denoting a generic posterior predictive distribution by $g(x)$ and its cumulative by $G(x)$, while the true value it ought to be compared to by y_{true}, the two performance measures takes the form

$$\text{logScore} = -g(y_{\text{true}}) \quad \text{and} \quad \text{CRPS} = \int_{-\infty}^{\infty}(G(x) - \mathbb{1}_{\{x>y_{\text{true}}\}}(x))^2\,\mathrm{d}x,, \tag{16}$$

where $\mathbb{1}_{\{\cdot\}}(x)$ denotes the indicator function. For both measures, smaller values reflect a better model. While the former rewards pointiness of the posterior predictive distribution by solely considering the point mass at y_{true}, the latter also considers the rest of the distribution, by punishing mass stronger the further away from y_{true} it is.

The above validation procedure enables us to readily compare our new suggested procedure to the references methods, and assess their strength and weaknesses as opposed to each other. However, it is difficult to judge the quality of the uncertainty assessment produced by a method using the performance measures in (16). To properly assess this,

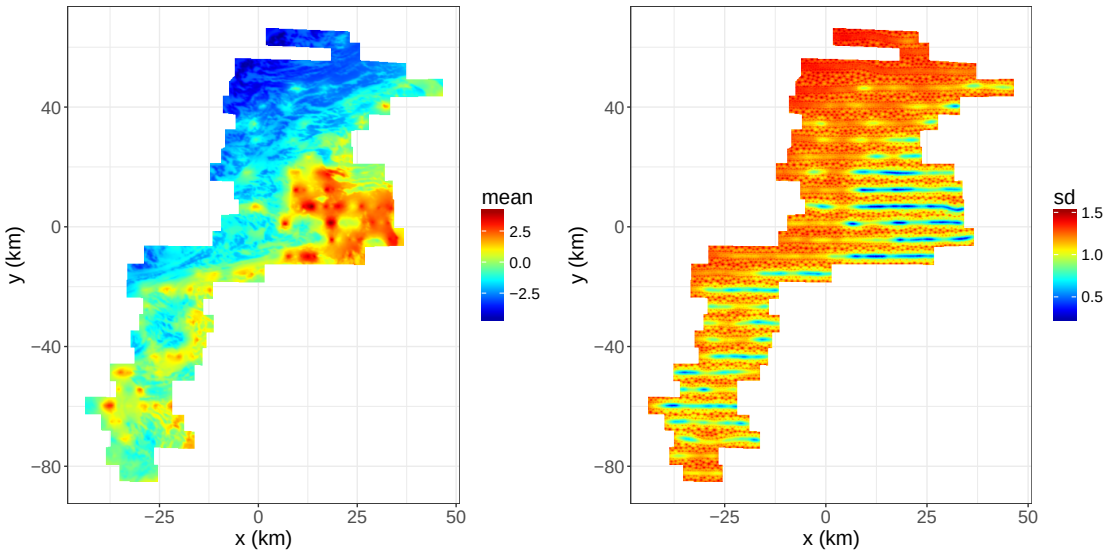

Figure 4. The mean and standard deviation of the latent field fitted for hooded seals with our procedure.

we suggest looking at the frequentist coverage of the posterior predictive distributions. That is, check how often the y_{true} lies within different credibility intervals, compared to their intended coverage – small uncertainty is of no value if it is not reflecting the true uncertainty of the model. To more closely mimic quantification of the complete whelping region, we suggest doing such an exercise on transect level.

4 Results

This section contains the results obtained when applying our procedure to the per-photo count data from the 2012 survey of the Greenland sea whelping region. As mentioned we model the hooded and harp seals separately. We compare our approach with the GAM-based procedure, both with a negative-binomial distribution for the counts and with the simpler Poisson distribution. As a baseline model we use a homogenous Poisson model with no covariates, spatial term, or other random effects. The models are compared through the aforementioned verification procedure.

4.1 Hooded seals

Within the flight transect sparsely covering the whelping region, a total of 777 hooded seal pups were counted. The blue dots in Figure 1 show how these are spread on the 2792 photos.

Figure 4 shows the mean and standard deviation of the latent field fitted using our procedure, as outlined in Section 3.2. As seen, the latent field captures the high intensity of hooded seal pups in the middle-eastern part of the whelping region. This area has a medium range ice-thickness, as seen from Figure 2. There is also an increased seal pup intensity further south, in particular closer to the open water. In the north, the intensity

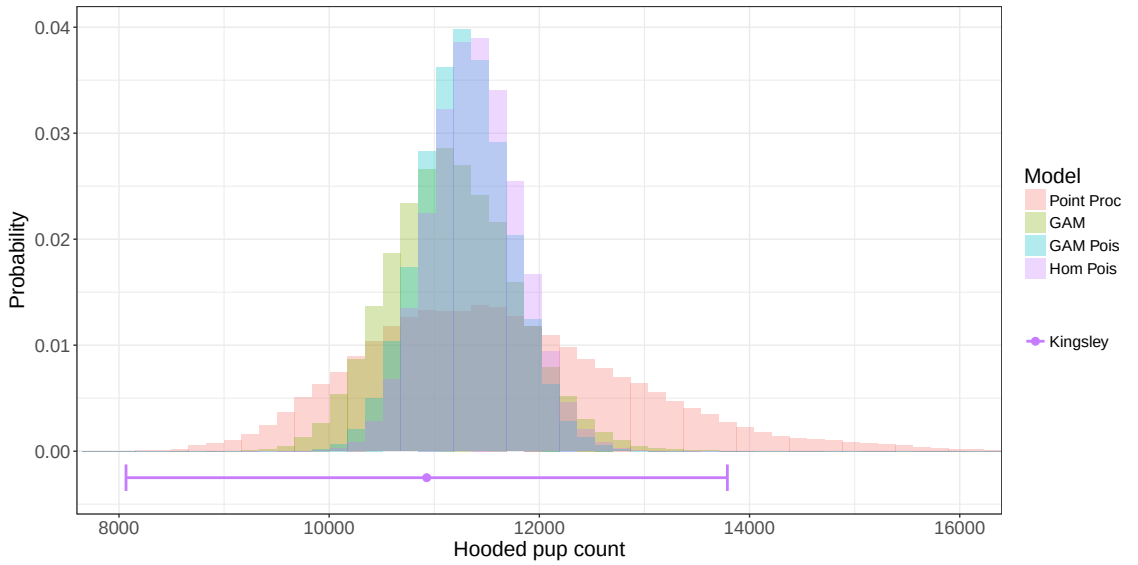

Figure 5. The posterior predictive distributions for the total area hooded pup counts for the three different competing models. For Kingsley's method, we show the point estimate +/- 2 times the estimated standard deviation, corresponding to an approximate 95% confidence interval under a normal distribution assumption.

is rather low. Note however, as seen from the sd-plot, that the uncertainty is rather large where the intensity is low, while it is smaller where the intensity is high. This means that one is fairly certain that there are *some* seal pups in areas where seal pups are observed, while there could very well be some seal pups even if none are observed in observation sites in the neighborhood. The range of the latent field, defined as the distance where the spatial correlation is approximately 0.1, has a posterior mean of 3.63km, that is about 2/3 of the distance between each transect. This means that the area of the latent field lying between transects essentially is determined by the neighboring transects. The fitted model gives the following posterior means for the intercept (α) and the fixed effects (β): $\text{mean}_\alpha = -1.37, \text{mean}_{\beta,q} = 9.07, \text{mean}_{\boldsymbol{\beta},s_1,s_2,s_{12}} = (0.07, -0.02, 3.63)$.

Figure 5 shows the posterior predictive distribution for the total area pup counts using our procedure, along with the corresponding results for the two GAM-based procedures and the homogenous Poisson model. A simple summary of Kingsley's method is also given for reference. Table 1 summarizes these distributions. As seen from the figure and read-off from the table, most of the mass coincide for the three methods, but our method is more conservative. The various point estimates are in the same ballpark, but with a tendency towards slightly higher estimates for our approach.

Table 2 shows the results from the validation scheme applied to the four methods we compare here, as outlined in Section 3.4. Starting with the photo level, the only significant result, defined as having non-overlapping 90% CI strictly below all others, is that our method is the best in terms of CRPS when using the random 10-fold CV scheme. It is also almost significant in terms of the logScore. When leaving out a full transect at a time, our method and the GAM version with the underlying negative binomial distribution

Table 1. Summary tables for the posterior predictive distributions for the total area counts of hooded seals.

	mean	median	mode	IQR	0.025-quantile	0.975-quantile
Point Proc	11649	11503	11472	1699	9472	14741
GAM	11178	11157	11093	807	10075	12395
GAM Pois	11296	11292	11093	572	10467	12147
Hom Pois	11391	11384	11377	572	10586	12238

Table 2. Validation results on respectively photo and transect level. Values in parenthesis are shows the lower and upper bounds of 90% bootstrapped confidence intervals for the performance measure. Cells shown in italics are the best (smallest) per column. Those which are significantly smaller than the others (defined as having non-overlapping confidence intervals) are also bolded.

HOODED SEALS: PHOTO LEVEL

	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
Point Proc	***0.18 (0.16, 0.19)***	*0.47 (0.44, 0.49)*	*0.22 (0.20, 0.25)*	0.54 (0.51, 0.57)
GAM	0.21 (0.19, 0.23)	0.51 (0.47, 0.53)	*0.22 (0.20, 0.24)*	*0.53 (0.50, 0.56)*
GAM Pois	0.22 (0.20, 0.24)	0.54 (0.51, 0.58)	0.24 (0.22, 0.26)	0.58 (0.54, 0.62)
Hom Pois	0.26 (0.24, 0.28)	0.69 (0.65, 0.73)	0.26 (0.24, 0.29)	0.70 (0.66, 0.74)

HOODED SEALS: AGGREGATE/TRANSECT LEVEL

	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
Point Proc	5.43 (4.04, 6.99)	3.68 (3.51, 3.86)	9.91 (5.99, 14.80)	*3.67 (3.26, 4.09)*
GAM	5.93 (4.95, 7.00)	3.79 (3.68, 3.91)	*9.37 (5.66, 13.63)*	3.68 (3.11, 4.27)
GAM Pois	5.90 (4.49, 7.42)	3.72 (3.50, 3.96)	10.14 (5.86, 15.09)	4.14 (3.32, 5.01)
Hom Pois	*4.83 (3.27, 6.66)*	*3.59 (3.30, 3.90)*	15.57 (11.77, 19.68)	6.71 (5.91, 7.52)

perform very similar, and somewhat better than the others.

Turning to the aggregate/transect level the baseline homogenous Poisson seems, perhaps somewhat surprisingly, to perform very well with the random leave-out scheme. A reasonable explanation for this is that since the photos are removed randomly, and what we measure here is the average of these, the baseline model, which essentially uses the overall average as its prediction, will naturally come close, and its low uncertainty will boost its performance in terms of these performance measures. The corresponding bad results when leaving out a full transect, shows that this model is generally way too simplistic.

In general, the validation results show that our method performs very similar to the original GAM-method, which are both clearly better than the other two alternatives. When it comes to predicting the seal pup count in individual randomly removed photos, our method is the clear winner. That is, our method is superior at predicting individual photos when it can use information about neighboring photo counts, which we think is an important property.

Although the point process and GAM methods can only be distinguished in terms of CRPS and logScore on photo level, their resulting posterior predictive distributions are quite different. To get further insight into this phenomenon, let us take a look at the posterior predictive distributions for the two methods per transect in the leave-out-transect setup. Figure 6 contains these distributions, plotted against the true transect count. As seen from the figure, the GAM method seems to have too narrow credibility intervals in terms of their frequentist coverage, while this matches better for the point process approach. In fact, the 90% CI covers the true count in $26/27 \approx 96\%$ of the transects for the point process approach, and only $18/27 \approx 67\%$ of the times for the GAM approach. The 50% CI is covered in respectively $16/27 \approx 59\%$ and $11/27 \approx 41\%$ of the transects for the point process and GAM approaches.

4.2 Harp seals

In total, 6034 harp seal pups where observed on the photos from the aerial photographic survey. As illustrated by the red dots in Figure 1, there are much larger packs of harp seal pups, than hooded seal pups. The maximal number of harp seals pups in a single photo is 160.

Figure 7 shows the mean and standard deviation of the fitted latent field using our procedure. Compared to the latent field for the hooded seal pups in Figure 4, the mean field has higher and steeper peaks. The places where the seal pups mainly appear are also similar to those for the hooded seal pups, except for some additional colonies in the north and north-west. Otherwise, the properties of the field are fairly similar to those for the hooded seal pups. The range of the latent field, has a posterior mean of 2.89, slightly smaller than for the Hooded seals. The fitted model gives the following posterior means for intercept (α) and fixed effects ($\boldsymbol{\beta}$): $\text{mean}_\alpha = -2.77, \text{mean}_{\boldsymbol{\beta},q} = 14.70, \text{mean}_{\boldsymbol{\beta},s_1,s_2,s_{12}} = (0.03, 0.01, -0.003)$. Note that the covariate effects are stronger for the harps than the hooded seals.

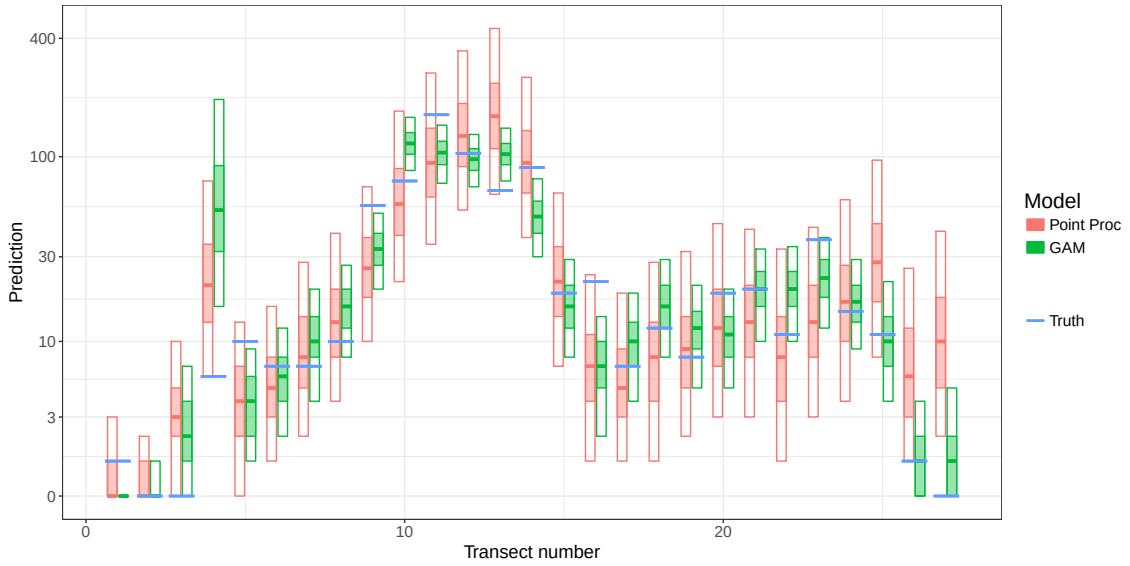

Figure 6. The posterior predictive distributions per transect for the point process and GAM methods plotted against the true count for hooded seals. For each method, the solid line shows the median, the light colored box shows the 50% CI, while the transparent box shows the 90% CI. The y-axis takes a $\log_{10}(x+2)$-scale to better show differences.

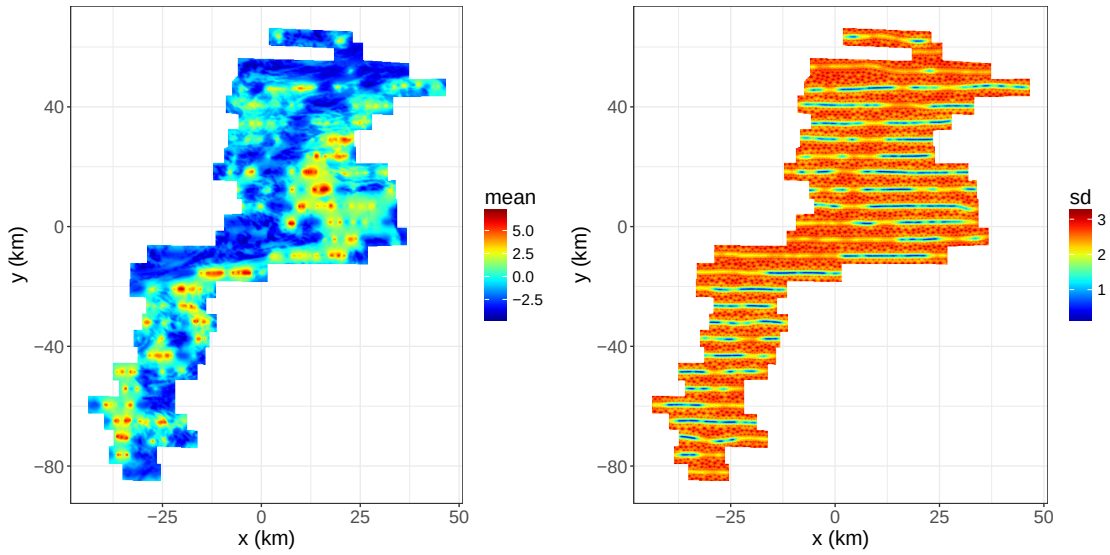

Figure 7. The mean and standard deviation of the fitted latent field for harps with our procedure.

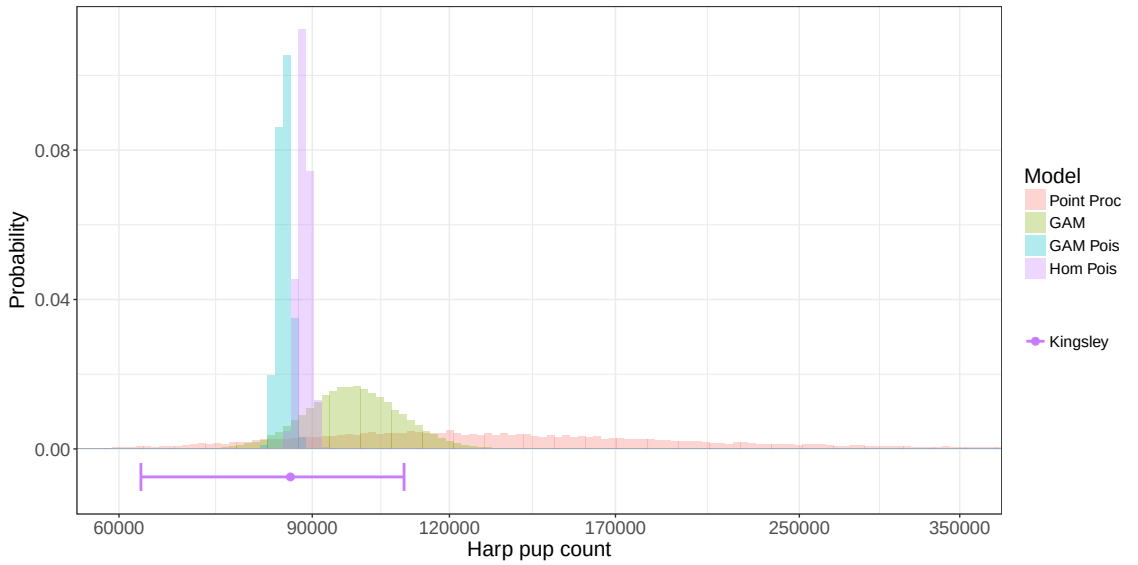

Figure 8. The posterior predictive distributions for the total area harp pup counts for the three different competing models. The x-axis is plotted on log-scale. For Kingsley's method, we show the point estimate +/- 2 times the estimated standard deviation, corresponding to an approximate 95% confidence interval under a normal distribution assumption.

Table 3. Summary tables for the posterior predictive distributions for the total area counts of harp seals.

	mean	median	mode	IQR	0.025-quantile	0.975-quantile
Point Proc	147919	127965	110996	72347	69267	357185
GAM	98617	98035	91876	12895	81023	119349
GAM Pois	84852	84852	84910	1536	82681	87094
Hom Pois	88424	88419	88389	1581	86149	90742

Figure 8 shows the posterior predictive distribution for the total number of harp seal pups in the whelping region using our procedure, along with the corresponding results for the two GAM procedures and the homogenous Poisson model. A simple summary of Kingsley's method is also given for reference. Table 3 summarizes these distributions. As seen from the figure and table, our point process is extremely vague about the harp seal pup count, essentially saying that the total number of seal pups could very well be above 250 000, but also less than 100 000. In contrast, the GAM method's upper tails ends at about 130 000 seal pups, while the two simplest methods has upper tails below 100 000.

Table 4 shows the results from the validation procedure for the Harp seal. On photo level the point process approach gives a significantly better CRPS and logScore with random 10-fold CV. Leaving out full transects gives no significantly best method although the GAM Poisson model tends to generally do well here.

As for the hooded seals, we take a closer look at the point process and GAM methods to better understand how well their different uncertainty estimates match the actual un-

Table 4. Validation results on respectively photo and transect level. Values in parenthesis show the lower and upper bounds of 90% bootstrapped confidence intervals for the performance measure. Cells shown in italics are the best (smallest) per column. Those which are significantly smaller than the others (defined as having non-overlapping confidence intervals) are also bolded.

HARP SEALS: PHOTO LEVEL

	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
Point Proc	***1.14 (1.01, 1.27)***	***0.95 (0.91, 1.00)***	1.96 (1.72, 2.20)	1.28 (1.22, 1.33)
GAM	1.78 (1.58, 2.00)	1.17 (1.11, 1.22)	*1.90 (1.67, 2.13)*	*1.27 (1.21, 1.33)*
GAM Pois	2.32 (2.10, 2.55)	2.09 (2.00, 2.17)	2.46 (2.22, 2.71)	2.17 (2.08, 2.26)
Hom Pois	2.68 (2.44, 2.94)	2.50 (2.45, 2.55)	2.69 (2.45, 2.95)	2.52 (2.46, 2.57)

HARP SEALS: AGGREGATE/TRANSECT LEVEL

	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
Point Proc	95.98 (51.20, 148.78)	6.57 (5.93, 7.33)	152.95 (111.93, 198.59)	*6.49 (5.98, 6.94)*
GAM	88.33 (70.94, 106.87)	*6.45 (6.31, 6.62)*	96.70 (61.05, 139.91)	7.00 (6.24, 7.76)
GAM Pois	*60.93 (42.57, 79.87)*	8.03 (6.83, 9.22)	*55.96 (39.62, 74.51)*	8.42 (7.34, 9.39)
Hom Pois	69.67 (45.62, 93.94)	7.69 (6.58, 8.73)	91.19 (61.09, 124.17)	7.96 (7.32, 8.57)

certainty. As expected, the GAM method has a much more narrow credibility interval which often does not cover the true seal pup count in the transect, while this matches the assigned credibility interval better for the point process approach. Out of the 27 transects, the 90% credibility intervals for the point process and the GAM approaches covers the true count in respectively $24/27 \approx 89\%$ and $18/27 \approx 67\%$ of the transects. The corresponding coverages for the 50% interval are $14/27 \approx 52\%$ and $11/27 \approx 41\%$. Thus, it is clear that from a frequentist perspective, the GAM method has too narrow credibility intervals, while our procedure get the level quite right.

5 Concluding remarks

We have presented a point process based procedure for estimating seal pup abundances based on observational data from an aerial photographic survey. Using the SPDE-INLA framework, we fit a Bayesian hierarhical model with Poisson counts based on a log-Gaussian Cox process formulation. As an additional contribution to seal pup abundance estimation, we adopt the use of satellite imagery as covariates in the modeling process, acting as a proxy for ice thickness. Our procedure is tested on data from a survey in the Greenland Sea in 2012, with both harp and hooded seal pup counts, and compared to three reference. Their performance is studied in a comprehensive (cross) validation study, were our procedure generally comes out best locally, while no method stands out as the

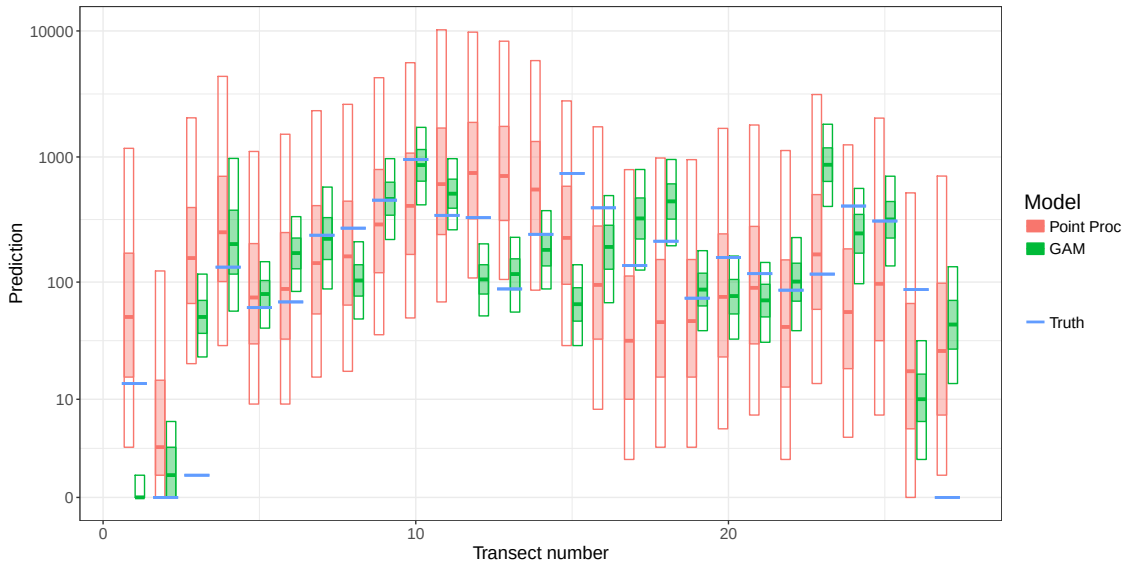

Figure 9. The posterior predictive distributions per transect for the point process and GAM methods plotted against the true count for harp seals. For each method, the solid line shows the median, the light colored box shows the 50% CI, while the transparent box shows the 90% CI. The y-axis takes a $\log_{10}(x+2)$-scale to better show differences.

best on a more global scale.

The most distinguishing character of our method is its larger uncertainty compared to the other methods. Our analysis suggests that our method has accurate frequentist coverage, while the reference methods are too optimistic in their uncertainty assessments. A possible reason for this is that we build a full Bayesian model from the ground with priors on parameters, while the other methods are merely pseudo-Bayesian procedures.

For the marine industry, this contribution gives both a new alternative for estimation of seal pup abundance, and a complete reference study showing the local and global properties of the different approaches. The frequentist coverage analysis also suggests to what extent the uncertainty output of the methods can be trusted. All of this is expected to contribute and guide the marine reseachers to better decision making in future, when monitoring the North Atlantic seal pup abundance.

References

Blangiardo, M. and Cameletti, M. (2015). *Spatial and Spatio-temporal Bayesian Models with R - INLA*. Wiley. 9

Cox, D. R. (1955). Some statistical methods connected with series of events. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 129–164. 7

Hastie, T. and Tibshirani, R. (1990). *Generalized additive models*. Wiley Online Library. 4

ICES (2014). Report of the ICES/NAFO working group on harp and hooded seals (WGHARP). Technical report, ICES CM 2014/ACOM:20, Quebec City, Quebec, Canada. 4

Kingsley, M., Stirling, I., and Calvert, W. (1985). The distribution and abundance of seals in the Canadian High Arctic, 1980–82. *Canadian Journal of Fisheries and Aquatic Sciences*, 42(6):1189–1210. 4, 12

Kovacs, K. M. and Lavigne, D. (1986). Cystophora cristata. *Mammalian Species*, (258):1–9. 4

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(4):423–498. 4, 8, 9, 10

Martins, T. G., Simpson, D., Lindgren, F., and Rue, H. (2013). Bayesian computing with INLA: New features. *Computational Statistics & Data Analysis*, 67:68–83. 8

Matheson, J. E. and Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management science*, 22(10):1087–1096. 14

Møller, J., Syversveen, A. R., and Waagepetersen, R. P. (1998). Log Gaussian Cox processes. *Scandinavian journal of statistics*, 25(3):451–482. 4, 8

Møller, J. and Waagepetersen, R. P. (2003). *Statistical inference and simulation for spatial point processes*. CRC Press. 4, 8

Øigård, T. A., Haug, T., and Nilssen, K. T. (2014a). Current status of hooded seals in the Greenland Sea. victims of climate change and predation? *Biological conservation*, 172:29–36. 4, 6

Øigård, T. A., Haug, T., and Nilssen, K. T. (2014b). From pup production to quotas: current status of harp seals in the Greenland Sea. *ICES Journal of Marine Science: Journal du Conseil*, 71(3):537–545. 4, 6

Øigård, T. A., Haug, T., Nilssen, K. T., and Salberg, A.-B. (2010). Estimation of pup production of hooded and harp seals in the Greenland Sea in 2007: Reducing uncertainty using generalized additive models. *Journal of Northwest Atlantic Fishery Science*, 42:103–123. 4, 13

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, 8, 10

Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P., and Lindgren, F. K. (2016). Bayesian computing with INLA: A review. *arXiv preprint arXiv:1604.00860*. 9

Salberg, A.-B., Haug, T., and Nilssen, K. T. (2008). Estimation of hooded seal (cystophora cristata) pup production in the Greenland Sea pack ice during the 2005 whelping season. *Polar Biology*, 31(7):867. 4, 12

Salberg, A.-B., Øigård, T. A., Stenson, G. B., Haug, T., and Nilssen, K. T. (2009). Estimation of seal pup production from aerial surveys using generalized additive models. *Canadian Journal of Fisheries and Aquatic Sciences*, 66(5):847–858. 4, 13, 14

Sergeant, D. E. (1974). A rediscovered whelping population of hooded seals cystophora cristata erxleben and its possible relationship to other populations. *Polarforschung*, 44(1):1–7. 4

Sergeant, D. E. (1991). Harp seals, man and ice. *Canadian Journal of Fisheries and Aquatic Sciences*, 114:1–153. 4

Simpson, D., Illian, J. B., Lindgren, F., Sørbye, S. H., and Rue, H. (2016). Going off grid: computationally efficient inference for log-Gaussian Cox processes. *Biometrika*, 103(1):49–70. 10

Simpson, D., Lindgren, F., and Rue, H. (2012). In order to make spatial statistics computationally feasible, we need to forget about the covariance function. *Environmetrics*, 23(1):65–74. 9