

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

You are reading the post-print version of this paper. Publisher's version is available here:<https://doi.org/10.1111/rssc.12397>

Martin Jullum

Norwegian Computing Center, Oslo, Norway.

E-mail: jullum@nr.no

Thordis Thorarinsdottir

Norwegian Computing Center, Oslo, Norway.

E-mail: thordis@nr.no

Fabian E. Bachl

School of Mathematics, University of Edinburgh, Edinburgh, Scotland.

E-mail: bachlfab@gmail.com

Summary. The Greenland Sea is an important breeding ground for harp and hooded seals. Estimates of the annual seal pup production are critical factors in the abundance estimation needed for management of the species. These estimates are usually based on counts from aerial photographic surveys. However, only a minor part of the whelping region can be photographed, due to its large extent. To estimate the total seal pup production, we propose a Bayesian hierarchical modeling 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 stochastic 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 results of the study indicate that our method improves local estimation performance, and that the increased prediction uncertainty of our method is required to obtain calibrated count predictions. This suggests that the sampling density of the survey design may not be sufficient to obtain reliable estimates of the seal pup production.

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 (Sergeant, 1974, 1991; Kovacs and Lavigne, 1986). 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. Both the

estimated total pup production *and* the associated uncertainty are then used as input in dynamic population models (Øigård et al., 2010).

The observational data consists of seal pup counts obtained by manual counting on photographs. The photographs stem 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 homogeneous, the negative binomial GAM approach and the Kingsley method yield similar estimates. However, the Kingsley method may possess a positive bias when the spatial distribution of the pups is clustered (Salberg et al., 2008; Øigård et al., 2010). Additionally, the GAM method produces much smaller uncertainty bounds than the homogeneous Kingsley 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 production 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 stochastic partial differential equation (SPDE) approach of the integrated nested Laplace approximation method (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. Although more traditional Markov Chain Monte Carlo (MCMC) methods in theory could be used to arrive at the same posterior, application of INLA allows results to be produced magnitudes faster, at a negligible cost in terms of accuracy.

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. This is important as the seal pups can only be observed on ice, with non-observable pups accounted for within the dynamic population model (Øigård et al., 2010). Compared to the other procedures our method gives larger uncertainties, especially for the harp seals. To validate these differences, we compare our proposed method with a number of reference methods in two cross-validation experiments, one where random sets of photos are removed and one where whole transects are removed from the data set prior to inference. Performance assessment based on proper scoring rules suggests our method performs best on a local level, and comparable on a more global scale. Further calibration assessment suggests that the larger uncertainty in our method is indeed more realistic.

The rest of the paper is organized as follows: Section 2 describes the survey 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. Relevant background related to point processes and aggregated point patterns are given in Section 3. Section 4 describes the details of our suggested modeling approach and three reference methods that we compare our method against. The validation schemes used to verify and compare the different approaches are also described. The results are presented in Section 5, including specifics of the model fitted with our procedure, and the validation and comparison results. The final Section 6 contains concluding remarks and pointers to future work. The online supplementary material contains a brief description of the INLA methodology and details about specification of the so-called mesh used by that methodology.

2. Data

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

2.1. Survey method

Before conducting the aerial photographic survey with the purpose of monitoring the seal pup production, the marine researchers typically perform a helicopter reconnaissance survey. This is done in order to locate the patches where the seals whelp for limiting 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 transects at a fixed distance that sparsely cover the survey area.

In this particular survey in March 2012, the airplane flew at an altitude of about 330m, and took a total of 2792 photos along 27 parallel transects, approximately 3Nm ($\approx 5.6\text{km}$) apart, with each photo covering $226 \times 346\text{m}$ of ground level. Due to fog, an exception was made for the two southernmost transects, which were flown at an altitude of 250m, with the photos covering $170 \times 260\text{m}$. Care is taken to record non-overlapping photos. Some overlaps appear in the data, but since the overlaps are so small and completely irrelevant for practical purposes, we assume non-overlap in the subsequent analysis. Along each transect, the cameras were turned on when the first seal was spotted from the airplane, and photos were taken continuously until the ice edge was reached on the eastern side, and until no seals were spotted for an extended period to the west. As a consequence of this survey setup, the whelping region is approximately defined as the union of the 1.5Nm ($\approx 2.8\text{km}$) bands around each transect. Thus, when estimating the total pup production in the whelping region, we only count predictions within this area. More details about the survey may be found in [Øigård et al. \(2014a,b\)](#).

2.2. Seal pup counts

Following the airborne survey, experienced marine researchers manually count the number of seal pups of each species in each photo[†]. Quality checks with multiple examinations are performed to limit the measurement error introduced in this step (Øigård et al., 2014a,b). The seal pup count data set used on the subsequent analysis contains the coordinates and extent of each photo, in addition to the number of seal pups of each species observed. The data are plotted in Figure 1 along with the transect locations and the extent of the whelping region. As seen from the figure, there tends to be more seal pups clustered towards the middle eastern boundary and southern corner of the whelping region. Comparatively more harp seal pups are observed than hooded seal pups and the spatial distribution of the former seems less homogeneous.




Fig. 1. Harp and hooded seal pup count data from the Greenland Sea 2012 survey, with east-west direction on the x -axis and north-south direction on the y -axis.

2.3. Satellite imagery

For whelping, the seals require large ice floes with access to the ocean for the adult seals to access food. Information regarding which areas are covered by ice floes and which areas are merely open water, is thus potentially highly relevant when estimating the seal pup production. In an attempt to account for this, we have collected high resolution satellite imagery (Modis) from the whelping region captured on the 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,

[†]The exact position of the seals are not recorded.

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

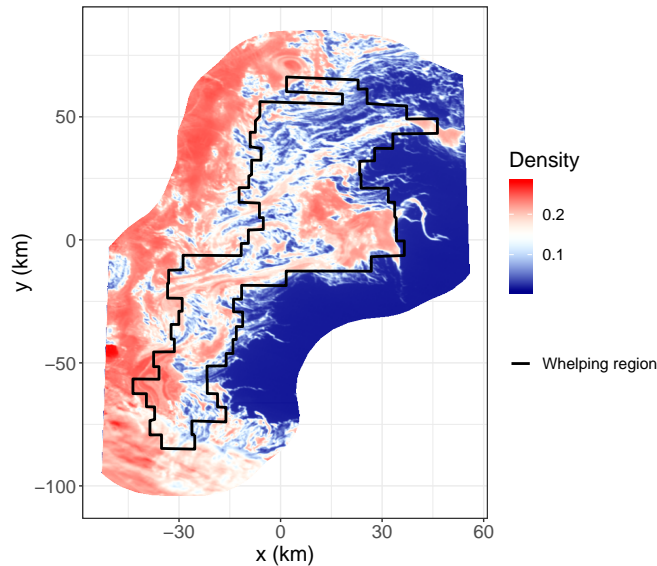


Fig. 2. Satellite data showing ice density used as a covariate in the model fitting process.

3. Background

In the present paper, we assume that the seal pup appearances can be thought of as a realization of a point process. The aim of the analysis is to estimate the total number of seal pups, $N(\Omega)$, in the whelping region $\Omega \subset \mathbb{R}^2$ (indicated in black in Figures 1 and 2).

3.1. Point processes

A spatial stochastic point process is a mathematical description of the random process through which points or locations are distributed in space. The collection of such observations is called a point pattern. In informal mathematical terms, a spatial point process Y is a random collection of points in a bounded observation region $\Omega \subset \mathbb{R}^2$ where both the number of points and their locations are random. The most fundamental type of point process is the *Poisson point process*. It may be specified by a deterministic intensity function $\lambda : \Omega \mapsto [0, \infty)$ which determines the (localized) expected density of points per unit area in Ω . The number of points, $N(B)$, of any Borel set $B \subseteq \Omega$ is Poisson distributed with mean $\mu(B) = \int_B \lambda(\mathbf{s}) d\mathbf{s}$, i.e. $N(B) \sim \text{Po}(\mu(B))$. Further, $N(B)$ is independent of $N(B^*)$ for any other non-overlapping Borel set $B^* \subseteq \Omega$ with $N(B \cup B^*) = N(B) + N(B^*)$.

The *Cox process* (also known as the doubly stochastic Poisson point process) introduced by [Cox \(1955\)](#) is a generalization of the Poisson point process where the intensity function λ is in 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 an underlying latent Gaussian random field Z , and given $\lambda = \exp(Z)$, Y is a Poisson point process with intensity λ . This type of model is quite flexible and useful for modeling a great variety of natural processes, in particular when only a single realization of the process is available (see e.g. Diggle, 2013; Illian et al., 2008; Møller and Waagepetersen, 2003). The LGCP model can describe various types of clustering in the point pattern through the positive semi-definite correlation function of the underlying Gaussian process (Møller et al., 1998).

Wolpert and Ickstadt (1998) consider a Cox process where the stochastic intensity function is given by $\lambda = \Gamma$ for a latent gamma random field Γ . As before, the conditional distribution of $N(B)$ given λ is a Poisson distribution. For this model, the marginal distribution of $N(B)$ is available in closed form and given by a negative binomial distribution (Matérn, 1971; Diggle and Milne, 1983).

3.2. Models for aggregated point patterns

As mentioned in Section 2.2, the exact positions of the counted seal pups are not available in the survey data. Instead, the seal pup counts are provided as aggregated counts per photo. Thus, our data is a partly observed point pattern aggregated to counts on an irregular lattice, as opposed to the actual point pattern. Methods that fit doubly stochastic Poisson process models to point pattern data, such as those proposed by Wolpert and Ickstadt (1998), Simpson et al. (2016) and Yuan et al. (2017), are thus not directly applicable. Rather, our setting can be viewed as an extension of the approach considered in Rue et al. (2009) where a fully observed point pattern is approximated by the corresponding counts on a *regular* lattice.

To this end, let $N(A_i)$ for $i = 1, \dots, n$ denote the number of seal pups in each of the $n = 2792$ photos with domains $A_i \subset \Omega$ and (potentially varying) areas $|A_i|$ for $i = 1, \dots, n$. Per properties of Poisson driven point processes, conditional on the intensity λ , being either deterministic (Poisson process) or random (Cox process), we have that $N(A_i)$ is Poisson distributed with parameter $\mu = \int_{A_i} \lambda(\mathbf{s}) \, d\mathbf{s}$, i.e.

$$N(A_i) | \lambda \sim \text{Po}\left(\mu = \int_{A_i} \lambda(\mathbf{s}) \, d\mathbf{s}\right), \quad i = 1, \dots, n. \quad (1)$$

Working under a Bayesian paradigm, we obtain an estimate of the total number of seal pups in Ω given by the posterior predictive distribution of $N(\Omega)$,

$$\begin{aligned} p(N(\Omega) | N(A_1), \dots, N(A_n)) &= \int p(N(\Omega), \lambda | N(A_1), \dots, N(A_n)) \, d\lambda \\ &= \int p(N(\Omega) | \lambda) p(\lambda | N(A_1), \dots, N(A_n)) \, d\lambda, \end{aligned} \quad (2)$$

where $N(\Omega) | \lambda \sim \text{Po}(\mu = \int_{\Omega} \lambda(\mathbf{s}) \, d\mathbf{s})$. The estimate for $N(\Omega)$ is thus given by a mixture of Poisson distributions rather than a single Poisson distribution under both deterministic and the random models for the intensity function λ . As the data are overdispersive, this may improve the fit of the Poisson model.

Table 1. Summary of observation models for seal pup counts per photograph and associated inference approaches. If the inference is performed based on an R-package, the package is listed in parenthesis. Here, s is a spatial index, λ_0, α and β are coefficients, \mathbf{x} is a vector of covariates, f denotes a latent spatial random effect, and S a deterministic spatial smoothing component. MPLE is an abbreviation for maximum penalized likelihood estimation.

Name	Observation model	Intensity	Inference approach
LGCP	Poisson	$\exp(\alpha + \beta^\top \mathbf{x}(s) + f(s))$	Bayesian with SPDE-INLA (INLA)
GAM Po	Poisson	$\exp(\alpha + \beta^\top \mathbf{x}(s) + S(s))$	MPLE with Bayesian estimator of uncertainty (<code>mgcv</code>)
GAM NB	Negative Binomial	$\exp(\alpha + \beta^\top \mathbf{x}(s) + S(s))$	MPLE with Bayesian estimator of uncertainty (<code>mgcv</code>)
Hom Po	Poisson	λ_0	Bayesian

The negative binomial distribution is a common choice for modeling random counts in applications with clustering. Due to the overdispersion in the data, [Salberg et al. \(2009\)](#) model the seal pup counts using a generalized additive model (GAM) based on a negative binomial likelihood. [Diggle and Milne \(1983\)](#) show that only two point process models yield a negative binomial marginal count distribution. One such model is the so-called compound Poisson process where each point of an underlying Poisson process is replaced by a random number of coincident points where the numbers are independent and identically distributed according to a logarithmic distribution. The second example is the Cox process with an intensity function given by a gamma random field considered by [Wolpert and Ickstadt \(1998\)](#). Similarly, if a Poisson process has a constant intensity and the inference is performed using a conjugate gamma prior distribution for λ , the resulting posterior predictive distribution in (2) is a negative binomial distribution, see Section 4.3 below.

4. Methods

We consider four different approaches to modeling the seal pup counts per photograph, see the summary in Table 1. Our new proposed method is based on the LGCP framework and accounts for potential clustering, or overdispersion, in the data through the underlying latent Gaussian random field. This approach is compared to the following reference approaches: A GAM approach previously proposed by [Salberg et al. \(2009\)](#) and [Øigård et al. \(2010\)](#) under both a Poisson and a negative binomial likelihood, and a homogeneous Poisson model. The methods and the associated inference approaches are described below. In addition, we also compare these spatial methods to a reference approach that estimates only the total number of seal pups in the region ([Kingsley et al., 1985](#)), see Section 4.3 below.

4.1. Stochastic intensity model

Our proposed approach (LGCP) is based on a Poisson point process with a stochastic intensity function (i.e. a Cox process). The stochastic intensity function takes the form

$$\lambda(\mathbf{s}) = \exp(Z_f(\mathbf{s})), \quad (3)$$

where Z_f is a Gaussian random field, or an LGCP with a conditional distribution for the aggregated counts on the form of (1). The continuous Gaussian random field takes the form

$$Z_f(\mathbf{s}) = \alpha + \boldsymbol{\beta}^\top \mathbf{x}(\mathbf{s}) + f(\mathbf{s}), \quad (4)$$

where α is an intercept term, $\boldsymbol{\beta}$ are regression coefficients for $\mathbf{x}(\mathbf{s}) = (q(\mathbf{s}), s_1, s_2, \sqrt{s_1^2 + s_2^2})^\top$, with $q(\mathbf{s})$ containing the ice density variable from the satellite imagery, while s_1, s_2 , and $\sqrt{s_1^2 + s_2^2}$ model linear spatial effects for $\mathbf{s} = (s_1, s_2)^\top \in \Omega$. Finally, $f(\mathbf{s})$ is a (non-linear) continuous Gaussian random field meant to model spatial dependence not captured by the covariates in \mathbf{x} . Specifically, $f(\mathbf{s})$ is given the Matérn covariance function of the form

$$\text{Cov}(f(\mathbf{s}), f(\mathbf{t})) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} (\kappa\|\mathbf{s} - \mathbf{t}\|)^\nu K_\nu(\kappa\|\mathbf{s} - \mathbf{t}\|), \quad (5)$$

for $\mathbf{s}, \mathbf{t} \in \Omega$ 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. Further, for identifiability of the intercept term in (4), we restrict $f(\mathbf{s})$ to integrate to zero over the modeling region.

Note that the model defined by (1), (3) and (4) is equivalent to the Poisson log-normal model of Christensen and Waagepetersen (2002) which is a special case of the spatial generalized linear mixed model framework proposed by Diggle et al. (1998). Christensen and Waagepetersen (2002) employ a Markov chain Monte Carlo (MCMC) algorithm for model inference using a slightly simpler covariance structure than the Matérn covariance function in (5) to analyze a data set where the observation sets $\{A_i\}_{i=1}^n$ are given by discs of fixed radius. Similar methods based on Poisson kriging are discussed in e.g. Bellier et al. (2010) and De Oliveira (2014) where more traditional geostatistical inference methods are employed.

4.1.1. Inference

To obtain an approximation to the posterior predictive distribution in (2), we apply the integrated nested Laplace approximation (INLA) of Rue et al. (2009) that allows for computationally feasible approximate Bayesian inference with *discrete* space latent Gaussian models, see Section 1 in the online supplementary material for a brief description. The stochastic partial differential equation (SPDE) approach of Lindgren et al. (2011) extends the INLA framework to also handle models with *continuous* latent fields as in (4). The SPDE approach is based on transforming the continuous latent field to a certain Gaussian Markov random field (GMRF), formulated through the solution of a SPDE. The key point is to approximate the continuous field $Z(\mathbf{s})$ by a field $Z_{\text{GMRF}}(\mathbf{s})$

living on a triangular mesh. For a triangular mesh with m triangle vertices, we write

$$Z_{\text{GMRF}}(\mathbf{s}) = \sum_{j=1}^m z_j \phi_j(\mathbf{s}), \quad (6)$$

where $\mathbf{z} = (z_1, \dots, z_m)^\top$ is a multivariate Gaussian random vector and $\{\phi_j(\mathbf{s})\}_{j=1}^m$ is a set of deterministic linearly independent basis functions which are piecewise linear between the vertices and chosen such that $\phi_j(\mathbf{s})$ is 1 at vertex j , and 0 at all other vertices. A consequence of the representation in (6) is that $Z_{\text{GMRF}}(\mathbf{s})$ is fully determined by \mathbf{z} . That is, $Z_{\text{GMRF}}(\mathbf{s})$ takes the value z_j at vertex j while its values inside the triangles are determined by linear interpolation.

Assume now that $Z(\mathbf{s})$ is equipped with the Matérn covariance function in (5). As this type of field is a solution to a certain SPDE, the precision matrix Q of \mathbf{z} takes an analytical form which can be approximated by a sparse matrix \tilde{Q} . Since $Z_{\text{GMRF}}(\mathbf{s})$ is completely determined by \mathbf{z} , this allows continuous field computations to be carried out approximately using the INLA implementation. Note that following certain guidelines for constructing the triangular mesh, the resulting approximation error 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 e.g. Blangiardo and Cameletti (2015) and Rue et al. (2016).

4.1.2. Model specification and fitting

The model parameters to be estimated are the regression parameters $(\alpha, \boldsymbol{\beta})$ in (4) and the hyperparameters (ν, κ, σ^2) of the Matérn covariance function in (5). To improve identifiability, we fix the Matérn smoothing parameter at $\nu = 2$, as is common when applying the INLA framework (see e.g. Blangiardo and Cameletti, 2015). The other hyperparameters related to the latent field are equipped with mesh dependent default priors specified by the INLA software. For our mesh these are $\theta_1 \sim \mathcal{N}(1.328, 10)$ and $\theta_2 \sim \mathcal{N}(-2.594, 10)$, where $\theta_1 = \log(\tau)$ and $\theta_2 = \log(\kappa)$, with $\sigma^2 = 1/(4\pi\kappa^2\tau^2)$. The intercept term α is assigned the improper prior $\mathcal{N}(0, \infty)$ while $\boldsymbol{\beta} \sim \mathcal{N}(0, 1000\mathcal{I}_4)$, where \mathcal{I}_4 denotes the 4×4 dimensional identity matrix. All these priors are non-informative and they seem to influence the final results to a very limited degree. The alternative use of penalized complexity (PC) priors (Simpson et al., 2017) resulted in very similar final results.

Using the triangular mesh described in Section 2 of the supplementary material, the INLA software produces posterior distributions for all individual hyperparameters, and enables sampling from the posterior of the complete latent field

$$p(Z | N(A_1), \dots, N(A_n)),$$

through the approximation of $Z(\mathbf{s})$ by $Z_{\text{GMRF}}(\mathbf{s})$, see Blangiardo and Cameletti (2015, Ch. 8.2). This allows us to use a Monte Carlo approximation for the integral in (2):

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

where \tilde{Z}_k is the k -th sample of the posterior latent field. We use $K = 10\,000$ in this and similar Monte Carlo integrations throughout the paper. Further, by the point process properties,

$$p(N(\Omega) | Z = \tilde{Z}_k) \sim \text{Po}\left(\mu = \int_{\Omega} \exp(\tilde{Z}_k(\mathbf{s})) \, d\mathbf{s}\right).$$

The integral $\int_{\Omega} \exp(\tilde{Z}_k(\mathbf{s})) \, d\mathbf{s}$ can be solved by e.g. a simple Riemann midpoint rule. This is achieved by dividing Ω into $J \approx 40\,000$ rectangles B_1, \dots, B_J centered in $\mathbf{s}_1, \dots, \mathbf{s}_J$, similar in size to the individual photos $\{A_i\}_{i=1}^n$, and using $\int_{\Omega} \exp(\tilde{Z}_k(\mathbf{s})) \, d\mathbf{s} \approx \sum_{j=1}^J \exp(\tilde{Z}_k(\mathbf{s}_j)) |B_j|$ where the values $\tilde{Z}_k(\mathbf{s}_j)$ are derived using (6). To find the J sets, the full modeling region is covered by a regular grid and the cells within the whelping region are selected. The final approximation to the posterior predictive distribution is thus

$$p(N(\Omega) | N(A_1), \dots, N(A_n)) \approx \frac{1}{K} \sum_{k=1}^K \text{Po}\left(\mu = \sum_{j=1}^J \exp(\tilde{Z}_k(\mathbf{s}_j)) |B_j|\right), \quad (7)$$

i.e. a Poisson mixture distribution.

4.2. Inhomogeneous intensity model

An alternative to the stochastic intensity function in (4) is a deterministic inhomogeneous intensity. With only one realization of the point pattern, it is generally not possible to distinguish between a doubly stochastic Poisson process and an inhomogeneous Poisson process (Møller and Waagepetersen, 2003). However, the underlying assumptions regarding the data structure and the components of the intensity model vary somewhat so that the resulting predictive distributions may differ. In particular, the LGCP assumes clustering through positive correlation between nearby locations in the underlying random field model while for an inhomogeneous Poisson process, each point is stochastically independent of all the other points in the process (see e.g. Møller and Waagepetersen, 2003).

The models we use here are motivated by the work of Salberg et al. (2009) who model the seal pup production by a generalized additive model (GAM) based on a negative binomial likelihood, except that they do not include covariate information. Apart from adding covariates, the below approach follows that of Salberg et al. (2009) and Øigård et al. (2010). Salberg et al. (2009) argue that a negative binomial likelihood should be used rather than a Poisson likelihood due to overdispersion in that data. However, as the inhomogeneous intensity and the Poisson mixture in the posterior predictive distribution may improve the fit of the Poisson model, we find it natural to include the inhomogeneous Poisson model in our list of models.

4.2.1. Poisson model

The inhomogeneous Poisson model takes the conditional Poisson form of (1) with intensity function $\lambda(\mathbf{s}) = \exp(\eta(\mathbf{s}))$ given by

$$\eta(\mathbf{s}) = \alpha + \boldsymbol{\beta}^\top \mathbf{x}(\mathbf{s}) + S(\mathbf{s}), \quad (8)$$

where $S(\cdot)$ is a spatial smoothing component given by a thin-plate smoothing regression spline (Wood, 2003), i.e. a smooth, nonlinear deterministic spatial effect. As noted by Wood (2017, Ch. 5.8) there is a certain duality between the smoothing component $S(\cdot)$ in (8) and the Gaussian random field $f(\cdot)$ in (4) in that $S(\cdot)$ can be interpreted as the mean of a latent random field. An assessment of uncertainty in the estimation of $S(\cdot)$ then corresponds to an assessment of the uncertainty in the mean of the random field.

To fit the model in (8), we follow Salberg et al. (2009), and rely on the `gam` function in the R-package `mgcv` (Wood, 2017). This function fits the model

$$N(A_i) \sim \text{Po}(\mu = |A_i| \exp(\eta(\mathbf{s}_i))), \quad \text{for } i = 1, \dots, n,$$

with the $|A_i|$ values as fixed offsets with overlapping cubic regressions on a set of artificial knots in space, and A_i centered in \mathbf{s}_i . An unbiased risk estimator (UBRE) is used to select the right amount of smoothing (Wahba, 1990; Wood, 2006).

The method returns maximum penalized likelihood estimates (MPLE) for the regression parameters α, β and parameters ν associated with the spatial smoothing component $S(\cdot)$, $\theta = (\alpha, \beta, \nu)$ (see e.g. Wood, 2017). To account for the uncertainty involved in the Poisson model, results of Silverman (1985) are used to obtain a posterior distribution for θ of the form

$$N(\theta, \text{Cov}(\theta)), \quad (9)$$

see also Wood (2006) and Salberg et al. (2009). We sample from (9) and approximate the integral over Ω by a simple Riemann midpoint rule using the J rectangles B_1, \dots, B_J centered in $\mathbf{s}_1, \dots, \mathbf{s}_J$ with $B_1 \cup \dots \cup B_J \approx \Omega$, as in (7). This gives the approximation

$$p_{\text{GAM}} \text{Po}(N(\Omega) | N(A_1), \dots, N(A_n)) \approx \frac{1}{K} \sum_{k=1}^K \text{Po}\left(\mu = \sum_{j=1}^J |B_j| \exp(\eta_{\tilde{\theta}_k}(\mathbf{s}_j))\right), \quad (10)$$

where $\eta_{\tilde{\theta}_k}(\mathbf{s}_j)$, is the value of (8) using θ -parameters corresponding to the k -th sample from (9).

4.2.2. Negative binomial model

An alternative to the inhomogeneous Poisson model is to consider the negative binomial model

$$N(A_i) | \lambda \sim \text{NegBin}\left(\mu = \int_{A_i} \lambda(\mathbf{s}) \, d\mathbf{s}, \tau\right), \quad (11)$$

with $\lambda(\mathbf{s}) = \exp(\eta(\mathbf{s}))$ where $\eta(\mathbf{s})$ is given in (8), and shape parameter τ . This model is fit in the same manner as the Poisson model above using

$$N(A_i) | \lambda \sim \text{NegBin}(\mu = |A_i| \exp(\eta(\mathbf{s}_i)), \tau), \quad \text{for } i = 1, \dots, n.$$

In contrast to the Poisson distribution, the negative binomial distribution is not closed under addition for different mean values. Thus, to arrive at a full posterior predictive distribution for the total seal pup count $N(\Omega)$ under the model in (11), we employ

a sampling procedure that relies on an underlying conditional independence condition stating that conditional on μ and τ , the point counts in disjoint sets are independent. Again, using the J disjoint sets B_1, \dots, B_J , we sample counts

$$N_k(B_j) | \lambda \sim \text{NegBin}(\mu = |B_j| \exp(\eta_{\hat{\theta}_k}(\mathbf{s}_j)), \tau = \hat{\tau}_j), \quad \text{for } j = 1, \dots, J; k = 1, \dots, K,$$

where \mathbf{s}_j is the center point of B_j , $\eta_{\hat{\theta}_k}(\mathbf{s}_j)$ is sampled from (9) as in (10), and $\hat{\tau}_j$ is the estimated shape parameter for B_j . The posterior predictive distribution for $N(\Omega)$ is then given by the empirical distribution function of $\{N_k(\Omega)\}_{k=1}^K$ where

$$N_k(\Omega) = \sum_{j=1}^J N_k(B_j), \quad \text{for } k = 1, \dots, K.$$

Both the sampling from the asymptotic normal distribution in (9) and the conditional independence assumption employed above, are also used by [Salberg et al. \(2009\)](#) to quantify the uncertainty around the total seal pup production.

4.3. Homogeneous intensity model

A simple reference model is a homogeneous Poisson model with a constant intensity. That is, we set $\lambda(\mathbf{s}) \equiv \lambda_0$ for a fixed scalar λ_0 and assume

$$N(A_i) | \lambda \sim \text{Po}(\mu = |A_i| \lambda_0), \quad \text{for } i = 1, \dots, n.$$

For a Bayesian inference, we equip λ_0 with a non-informative conjugate gamma prior,

$$\lambda_0 \sim \Gamma(a_0 = 10, b_0 = 10),$$

for both seal types, where a_0 and b_0 are the shape and rate parameters, respectively. This results in the following posterior distribution,

$$\lambda_0 | N(A_1), \dots, N(A_n) \sim \Gamma\left(a = a_0 + \sum_{i=1}^n N(A_i), b = b_0 + \sum_{i=1}^n |A_i|\right).$$

Moreover, the posterior predictive distribution in (2) takes the form of a negative binomial distribution and is available in a closed form

$$\begin{aligned} p_{\text{Hom Po}}(N(\Omega) | N(A_1), \dots, N(A_n)) &= \int \text{Po}(|\Omega| \lambda_0) \Gamma(a, b) d\lambda_0 \\ &= \text{NegBin}(\mu = |\Omega| a/b, \tau = a). \end{aligned} \quad (12)$$

where a and b are specified in the posterior distribution for λ_0 above.

The motivation for the homogeneous Poisson model is not only that it is a simple special case of our proposed LGCP model, but also that it is similar to the traditional approach to estimate seal pup production based on aerial photographic transect surveys, often referred to as Kingsley's method ([Kingsley et al., 1985](#)). Kingsley's method is fundamentally simple: For each transect T_1, \dots, T_{27} covering the space $A_{T_k} = \bigcup_{\{A_i \subset T_k\}} A_i$,

compute the seal pup count $N_{T_k} = \sum_{A_i \in T_k} N(A_i)$. Then, the estimate of the total number of seal pups is

$$\hat{N}(\Omega) = \frac{|A_\Omega|}{\sum_{k=1}^{27} |A_{T_k}|} \sum_{k=1}^{27} N_{T_k}. \quad (13)$$

Kingsley et al. (1985) also provide an estimate of the variance related to the seal pup production estimate, based on serial differences between the transects. Salberg et al. (2008) later provided a modification to this variance estimate, which we have used here. Since this method is not spatial in nature and only provides an estimate for the total number of seal pups in the study area, it is difficult to properly compare it against the remaining spatial procedures in a cross-validation set-up where we assess the predictive performance at photo-level. We will therefore not perform validation tests, as described in Section 4.4, for this method.

4.4. Verification

We compare the various modeling approaches using a cross-validation scheme where we rely on two procedures for subsetting the data. The first procedure 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 second procedure 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 procedures, we fit the competing models 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 photos (corresponding to the full transect for the latter procedure).

We compare the predictive performance of the various modeling approaches using two performance measures: The logarithmic score (Good, 1952) and the continuous ranked probability score (CRPS) (Matheson and Winkler, 1976), which both are proper scoring rules that assess full predictive distributions (Gneiting and Raftery, 2007). Denoting a generic posterior predictive distribution by $g(x)$, the corresponding cumulative distribution function by $G(x)$, and the observed count by y_{true} , the two performance measures takes the form

$$\begin{aligned} \text{logScore}(g, y_{\text{true}}) &= -\log(g(y_{\text{true}})), \\ \text{CRPS}(G, y_{\text{true}}) &= \int_{-\infty}^{\infty} (G(x) - \mathbb{1}_{\{x > y_{\text{true}}\}}(x))^2 dx, \end{aligned} \quad (14)$$

where $\mathbb{1}_{\{ \cdot \}}(x)$ denotes the indicator function. For both measures, smaller values reflect a better model.

Both the logarithmic score and the CRPS are optimized in expectation when the true data distribution is issued as the forecast (Gneiting and Raftery, 2007). However, in real-life situations, average scores are often associated with high uncertainty (Thorarinsdottir and Schuhen, 2018). We thus apply two scores that penalize prediction errors in slightly different manners, see Figure 3. In particular, the logarithmic score is more sensitive to outliers, cf. the left plot in Figure 3. For every cross-validation schemes, we present mean




Fig. 3. Illustration of the logScore (solid lines) and CRPS (dashed lines) functions for different values of y_{true} , with two different predictive distributions (indicated in gray), $\text{Po}(\mu = 3)$ (left plot) and $\text{Po}(\mu = 20)$ (right plot).

scores, where the mean is computed over the different folds. To indicate the variation in the scores, we compute 90% bootstrapped confidence intervals (CI) of those means by repeated re-sampling (10 000 times, with replacement) of the fold scores.

We further assess the calibration, or the prediction uncertainty, of the methods by assessing the coverage of the posterior predictive distributions. That is, we check how often y_{true} lies within different credibility intervals, compared to their intended coverage – small prediction uncertainty is of no value if it is not reflecting the true variability of the data. For a calibrated forecast, an event predicted with probability p should be realized with the same frequency in the observed data. In order to mimic quantification of the prediction uncertainty for the complete whelping region as closely as possible, we perform this exercise on the transect level.

5. Results

Here, we present the results obtained when applying the various spatial models/estimation methods in Table 1 and Kingsley’s method to per-photo count data from the 2012 survey of the Greenland sea whelping region. We model hooded and harp seals separately as their occurrences are expected to be independent conditional on the covariate information. Specifically, we compare our LGCP approach estimated using SPDE-INLA with the GAM-based procedure, both with a negative-binomial distribution for the counts (GAM NB) and with the simpler Poisson distribution (GAM Po). As baseline models we use a homogeneous Poisson model with no covariates, spatial term, or other random effects (Hom Po), in addition to Kingsley’s method.

5.1. 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




Fig. 4. The mean and standard deviation of the random field $Z_f(s)$ in (4), fitted for hooded seals with our LGCP approach.

2792 photos, with between 0 and 12 pups per photo.

Figure 4 shows the mean and standard deviation of the random field fitted using our LGCP procedure outlined in Section 4.1. 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, cf. Figure 2. There is also an increased seal pup intensity further south, in particular closer to the open water. Generally speaking, as seen from the standard deviation plot, the uncertainty is rather large where the intensity is low, while it is smaller where the intensity is high. This means that apart from the north, one is fairly certain that there are *some* seal pups in areas where seal pups are observed nearby, while there could very well exist seal pups in locations where none are observed at the neighboring observation sites. As a result of very few seals being observed in the north, the mean intensity here is so low that it is unlikely that a significant amount of seals have settled there.

The *range* of the latent field, defined as the distance at which the spatial correlation is approximately 0.1, has a posterior mean of 3.63 km, or about 2/3 of the distance between two transects. This means that in an area lying between two transects, the latent field is essentially determined by the two neighboring transects. Further, the fitted model gives the following posterior means for the intercept (α) and the fixed effects (β): $\text{mean}_{\alpha} = -1.37$, $\text{mean}_{\beta,q} = 9.07$, $\text{mean}_{\beta,s_1,s_2,s_{12}} = (0.07, -0.05, -0.02)$. Similarly the GAM NB approach has the following coefficient estimates: $\hat{\alpha} = -2.58$, $\hat{\beta}_q = 9.59$, and $\hat{\beta}_{s_1,s_2,s_{12}} = (0.06, -0.02, 0.03)$.

Figure 5 shows the posterior predictive distribution for the total pup count in the whelping region using our LGCP procedure, along with the corresponding results for the two GAM-based procedures (negative binomial and Poisson response) and the homogeneous Poisson model. A simple summary of Kingsley's method is also given for

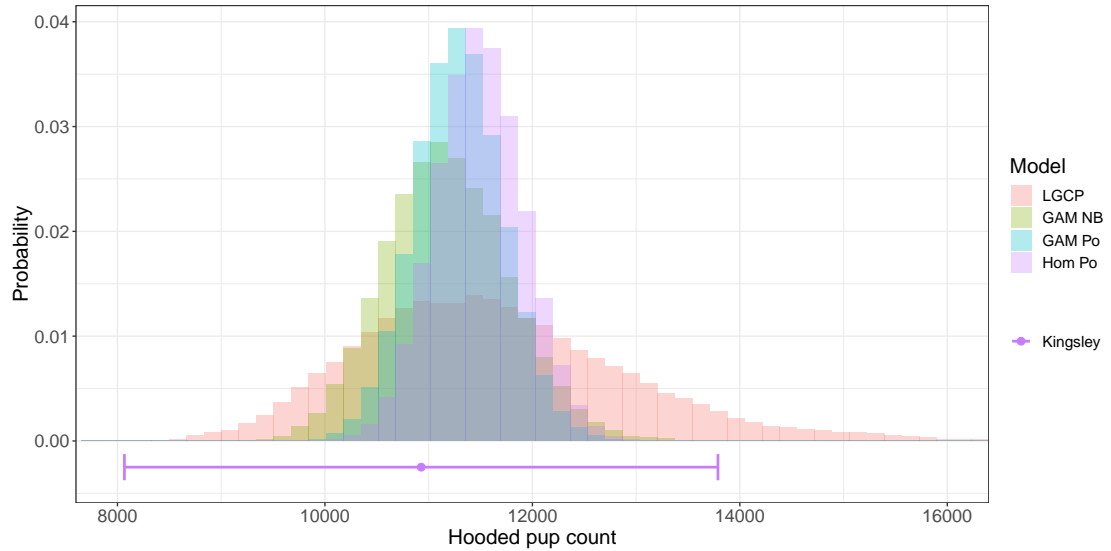


Fig. 5. The posterior predictive distributions for the total hooded pup counts in the whelping region for the five competing models. For Kingsley's method, we show the point estimate ± 2 times the estimated standard deviation, corresponding to an approximate 95% confidence interval under a normal distribution assumption.

Table 2. Summary table for the predictive distributions of the total count of hooded seal pups in the whelping region.

	mean	median	mode	IQR	0.025-quantile	0.975-quantile
LGCP	11649	11503	11472	1699	9472	14741
GAM NB	11178	11157	11093	807	10075	12395
GAM Po	11296	11292	11093	572	10467	12147
Hom Po	11494	11489	11479	571	10678	12338
Kingsley	10928	10928	10928	1969	8067	13789

Table 3. Validation results on photo level (one prediction per photo) and transect level (one prediction per transect), respectively. Lower and upper bounds of 90% bootstrapped confidence intervals for the scores are shown in parenthesis. 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 in bold.

HOODED SEALS: PHOTO LEVEL				
	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
LGCP	<i>0.18 (0.16, 0.19)</i>	<i>0.47 (0.44, 0.49)</i>	<i>0.22 (0.20, 0.25)</i>	0.54 (0.51, 0.57)
GAM NB	0.21 (0.19, 0.23)	0.51 (0.47, 0.53)	<i>0.22 (0.20, 0.24)</i>	<i>0.53 (0.50, 0.56)</i>
GAM Po	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 Po	0.26 (0.24, 0.28)	0.77 (0.71, 0.84)	0.26 (0.24, 0.29)	0.78 (0.72, 0.85)

HOODED SEALS: AGGREGATE/TRANSECT LEVEL				
	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
LGCP	5.43 (4.04, 6.99)	3.68 (3.51, 3.86)	9.91 (5.99, 14.80)	<i>3.67 (3.26, 4.09)</i>
GAM NB	5.93 (4.95, 7.00)	3.79 (3.68, 3.91)	<i>9.37 (5.66, 13.63)</i>	3.68 (3.11, 4.27)
GAM Po	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 Po	<i>4.89 (2.21, 11.06)</i>	<i>3.58 (3.14, 4.65)</i>	20.70 (1.29, 55.68)	12.81 (2.46, 37.28)

reference. Table 2 summarizes the predictive distributions. While the confidence bands for the five methods largely overlap, Kingsley’s method yields the lowest point estimates (mean, median and mode) and the LGCP method the highest estimates, closely followed by the Hom Po approach. The Hom Po approach and the GAM Po methods have the lowest prediction uncertainty. Here, our LGCP method has an interquartile range approximately twice as large as the GAM NB approach, and three times larger than GAM Po and Hom Po approaches. For all the spatial methods, the mean predictor is slightly higher than the median predictor, indicating a small degree of skewness with a heavier upper tail. This effect is most pronounced for the LGCP approach.

Table 3 shows the results from the validation scheme applied to the four methods we compare here, as outlined in Section 4.4. At the photo level, we issue a prediction for the pup count per photo, for either 10% of the photos or all photos in a single transect at a time. Here, our LGCP method yields very good results, in particular for the random 10-fold cross-validation where observations are generally available in the neighborhood of the prediction locations. It is significantly the best method for this setting under the CRPS, defined as having non-overlapping 90% CI strictly below all others, and almost significant in terms of the logScore. When leaving out a full transect at a time, our LGCP method and GAM NB method perform very similar, and somewhat better than the others.

At the aggregate/transect level we issue a joint prediction for the total pup count per 10% of the photos or per transect, respectively. Here, the baseline homogeneous Poisson model performs very well with the random leave-out scheme, indicating that the data may be close to homogeneous across the photos. However, this does not transfer to the second set-up where we leave out one transect at a time in which case our LGCP




Fig. 6. The predictive distributions per transect for the LGCP and GAM NB 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 has a $\log_{10}(x + 2)$ -scale to better show differences.

method and the GAM NB method again perform well. As seen in Figure 1, the data do not appear homogeneous across the transects. Note that these reported scores are averages over relatively few predictions, only 10 distinct predictions for the random 10-fold cross-validation study and 27 distinct predictions for the leave-out-transect setup. The score values are thus associated with a large degree of uncertainty. The uncertainty is particularly high for the homogeneous Poisson model that, due to the assumption of homogeneity, is especially sensitive to inhomogeneities across the different cross-validation sets.

Based on the results in Table 3, our LGCP method performs somewhat similar to the GAM NB method and these are both clearly superior to the other two alternatives; the two methods can only be distinguished in terms of CRPS on the photo level. Despite this, their resulting posterior predictive distributions are quite different. To get further insight into this phenomenon, Figure 6 shows the posterior predictive distributions for the two methods per transect in the leave-out-transect setup, plotted against the true transect counts. As seen from the figure, the GAM NB method seems to have too narrow credibility intervals, while the LGCP approach appears more calibrated. In fact, the 90% CI covers the true count in $26/27 \approx 96\%$ of the transects for the LGCP approach, and only $18/27 \approx 67\%$ of the transects for the GAM NB approach. The 50% CI is covered in respectively $16/27 \approx 59\%$ and $11/27 \approx 41\%$ of the transects for the LGCP and GAM NB approaches. Note that when looking at the MAE (mean absolute error) of the median count estimate for the two methods, the GAM NB method achieves MAE of 12.4, while the LGCP method is slightly worse with MAE of 13.8. Thus, the GAM NB method seems to do somewhat better as a point estimator, while the LGCP is better at




Fig. 7. The mean and standard deviation of the random field $Z_f(s)$ in (1) fitted for harp seals with our LGCP approach.

estimating the uncertainty.

5.2. Harp seals

A total of 6034 harp seal pups were observed on the 2792 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, indicating a higher degree of inhomogeneity. Here, the pup count per photo ranges from 0 to 160.

Figure 7 shows the mean and standard deviation of the fitted random field $Z_f(s)$ using our LGCP procedure. Compared to the latent field for the hooded seal pups in Figure 4, the mean field here has a much higher degree of spatial variation with higher and steeper peaks. However, the locations where the seal pups mainly appear are similar to those for the hooded seal pups, except for some additional colonies in the north and north-west of the region. Otherwise, the properties of the two fields are fairly similar. The range of the latent field has a posterior mean of 2.89 km, a slightly smaller value than for the hooded seals which corresponds to roughly half the distance between two transects. The fitted model gives the following posterior means for intercept (α) and fixed effects (β): $\text{mean}_\alpha = -2.77$, $\text{mean}_{\beta,q} = 14.70$, $\text{mean}_{\beta,s_1,s_2,s_{12}} = (0.03, 0.01, -0.003)$. Note that the covariate effects are stronger for the harps than the hooded seals. This is natural as there are many more observed harps than hooded seals. The GAM NB approach gives the following coefficient estimates: $\hat{\alpha} = 2.70$, $\hat{\beta}_q = 19.05$, and $\hat{\beta}_{s_1,s_2,s_{12}} = (0.02, 0.04, -0.1)$.

Figure 8 shows the posterior predictive distribution for the total number of harp seal pups in the whelping region using our LGCP procedure, along with the corresponding results for the two GAM procedures (GAM NB and GAM Po) and the homogeneous Poisson model (Hom Po). A simple summary of Kingsley's method is also given for




Fig. 8. The predictive distributions for the total harp pup counts in the whelping region for the five different approaches. 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 4. Summary table for the predictive distributions for the total number of harp seal pups in the whelping region.

	mean	median	mode	IQR	0.025-quantile	0.975-quantile
LGCP	147919	127965	110996	72347	69267	357185
GAM NB	98617	98035	91876	12895	81023	119349
GAM Po	84852	84852	84910	1536	82681	87094
Hom Po	88272	88267	88257	1583	85986	90587
Kingsley	85968	85968	85968	15926	62829	109107

reference. Table 4 summarizes these distributions. The LGCP prediction of the total harp seal pup count is highly uncertain, 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 NB method’s upper tail ends at about 130 000 seal pups, while the GAM Po and Hom Po methods agree that there are between 80 000 and 90 000 harp seal pups within the whelping region. Kingsley’s method yields point estimates similar to GAM Po and Hom Po while its uncertainty (IQR) is tenfold greater, or similar to that for GAM NB. Thus, there are significant differences both between the centrality and width of the different methods’ posterior predictive distributions. The predictive distributions for GAM Po and Hom Po are essentially symmetric while the one for GAM NB is minimally skewed with a heavier upper tail. The LGCP approach, however, yields a severely skewed predictive distribution with a predictive mean that is 15% larger than the predictive median due to the large uncertainty in the estimation of the random field Z_f , cf. Figure 7.

To further investigate the differences between the four spatial predictions, we have

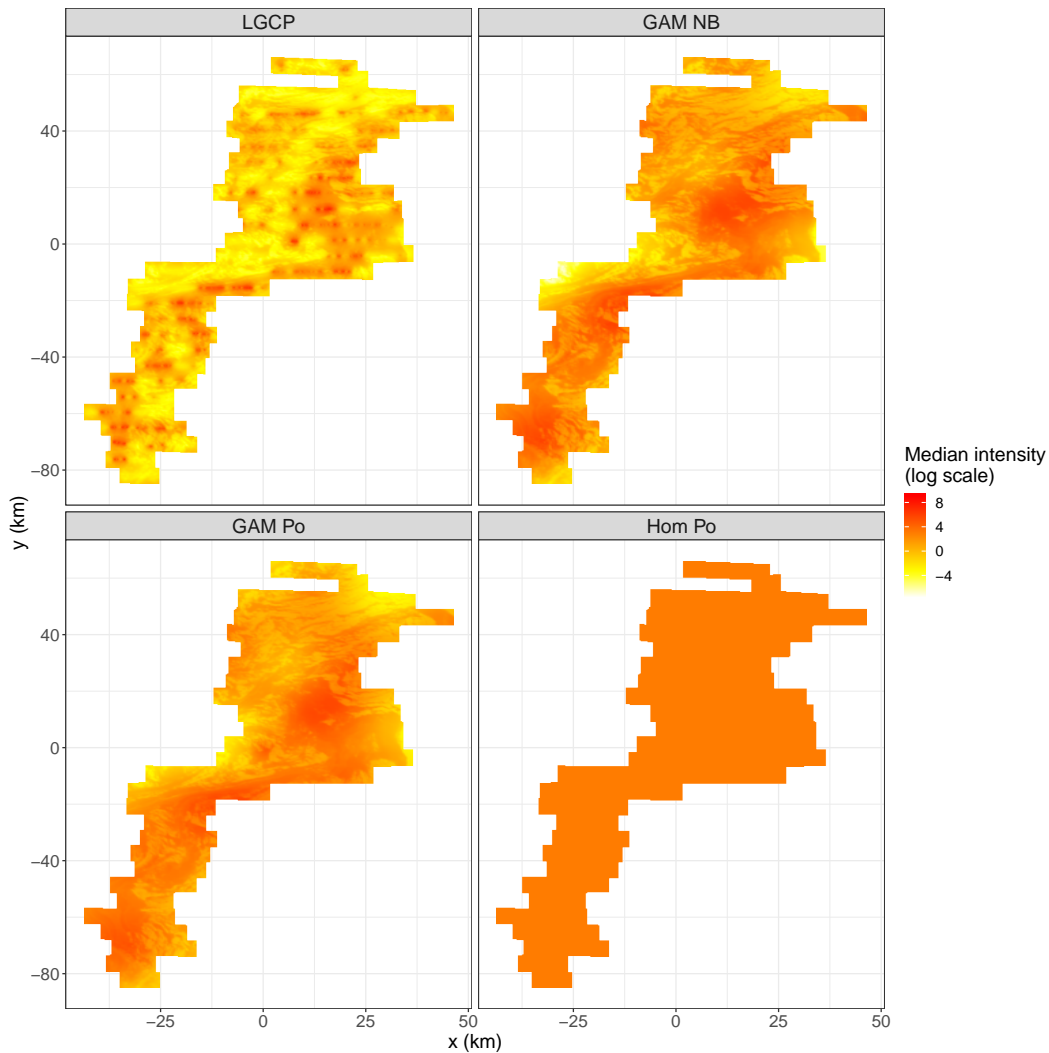


Fig. 9. The estimated predicted median intensity field (plotted on a logarithmic scale) for all four spatial methods fitted for harp seals.

plotted the predictive median intensity fields for the four methods in Figure 9, cf. the expressions for the intensity fields in the third column in Table 1. The regression coefficient for the ice density covariate q shown in Figure 2 is somewhat higher for the GAM approaches than for the LGCP approach (posterior mean of 19.05 for GAM NB compared to 14.70 for LGCP), resulting in fairly smooth median intensity fields that reflect the spatial structure of the ice density field. The median LGCP intensity field, however, has stronger inhomogeneities across space and appears more strongly influenced by the data density shown in Figure 1 with noticeable peaks in locations with higher observed seal pup density. Overall, however, the median intensity is higher for the GAM methods than for the LGCP. The reason that LGCP still produces higher predictions (cf. Table




Fig. 10. The posterior predictive distributions per transect for the LGCP and GAM NB methods plotted against the true count for harp seal pups. 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.

4) is partly due to the high peaks, but mainly due to the high degree of uncertainty. The posterior of $Z(\mathbf{s})$ is more or less symmetric with a high degree of uncertainty, such that large values are sampled quite frequently. These large values boosts the intensity $\lambda(\mathbf{s}) = \exp(Z(\mathbf{s}))$ considerably, resulting in large predictions. The lower uncertainty of the GAM methods does not have the same effect. Further, note that although the two GAM models have very similar posterior median intensities, their posterior predictive distributions are quite different, cf. Figure 8 and Table 4.

As for the hooded seals, we take a closer look at the LGCP and GAM NB methods to better understand how well their different estimates of the prediction uncertainty match the actual uncertainty, see Figure 10. As expected, the GAM NB method has a much narrower credibility intervals which too often fail to cover the true seal pup count in the transect, while our LGCP method shows much better calibration. Out of the 27 transects, the 90% credibility intervals for the LGCP and GAM NB approaches covers the true count in respectively $24/27 \approx 89\%$ and $18/27 \approx 67\%$ of the transects. The corresponding coverage for the 50% interval are $14/27 \approx 52\%$ and $11/27 \approx 41\%$, respectively. Thus, it is clear that the GAM NB method is underdispersive, while our procedure seems well calibrated. On the other hand, as for the hooded seals, the posterior median of the GAM NB method does better as a point estimator in terms of MAE (130.2) than our LGCP method (179.1).

Table 5 shows the results from the validation procedure for the harp seals, yielding similar model rankings as for the hooded seals. On photo level, the LGCP approach gives a significantly better CRPS and logScore under random 10-fold cross-validation. Leaving out full transects gives no significantly best method although the GAM Po

Table 5. Validation results on photo level (one prediction per photo) and transect level (one prediction per transect), respectively. Lower and upper bounds of 90% bootstrapped confidence intervals for the scores are shown in parenthesis. 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 in bold.

HARP SEALS: PHOTO LEVEL				
	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
LGCP	<i>1.14 (1.01, 1.27)</i>	<i>0.95 (0.91, 1.00)</i>	1.96 (1.72, 2.20)	1.28 (1.22, 1.33)
GAM NB	1.78 (1.58, 2.00)	1.17 (1.11, 1.22)	<i>1.90 (1.67, 2.13)</i>	<i>1.27 (1.21, 1.33)</i>
GAM Po	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 Po	2.64 (2.40, 2.90)	3.47 (3.28, 3.67)	2.66 (2.42, 2.92)	3.49 (3.30, 3.69)

HARP SEALS: AGGREGATE/TRANSECT LEVEL				
	Random 10-fold CV		Leave-out full transect	
	CRPS	logScore	CRPS	logScore
LGCP	95.98 (51.20, 148.78)	6.57 (5.93, 7.33)	152.95 (111.93, 198.59)	<i>6.49 (5.98, 6.94)</i>
GAM NB	88.33 (70.94, 106.87)	<i>6.45 (6.31, 6.62)</i>	96.70 (61.05, 139.91)	7.00 (6.24, 7.76)
GAM Po	<i>60.93 (42.57, 79.87)</i>	8.03 (6.83, 9.22)	<i>55.96 (39.62, 74.51)</i>	8.42 (7.34, 9.39)
Hom Po	102.62 (8.87, 296.09)	62.92 (4.31, 300.60)	149.38 (7.70, 454.53)	82.63 (4.08, 251.00)

method tends to generally do well here. However, as before, these average scores are associated with a very high degree of uncertainty so that a ranking of the methods based on these results is not advisable.

6. Conclusions and discussion

We have presented a point process based approach to estimate seal pup production based on observational data from an aerial photographic survey. Using the SPDE-INLA framework, we fit a Bayesian hierarchical model with Poisson counts following a log-Gaussian Cox process (LGCP) model formulation. As an additional contribution to seal pup production estimation, we adopt the use of satellite imagery as covariates in the modeling process, to act as a proxy for ice thickness. The approach is applied to 2012 survey data from the Greenland Sea, with both harp and hooded seal pup counts, and compared to several reference methods that can be associated with non-homogeneous or homogeneous point process formulations rather than the doubly stochastic setting of the Cox process.

The competing methods are compared in two cross-validation studies. The proposed LGCP approach generally performs best locally, while no method stands out as the best on a more regional scale. However, this lack of discrimination in the comparison at the regional scale is not surprising given the relatively small size of our data set resulting in large uncertainties in the scores, see e.g. the discussion and examples in [Thorarinsdottir and Schuhen \(2018\)](#). The most distinguishing characters of the LGCP method are higher count predictions and a large prediction uncertainty compared to

the other methods. Our analysis suggests that the wide uncertainty bounds are indeed necessary to issue calibrated predictions. This further suggests that the amount of data collected in the aerial photographic survey might not be sufficient to obtain reliable estimates of the total seal pup production in the whelping region.

The results show that the model uncertainty directly affects the seal pup count predictions to a large degree. This exemplifies the importance of a proper modeling of all involved uncertainties. In particular, a comparison of the two GAM approaches for the harp seals shows that even if the posterior median intensity fields for the two approaches are very similar, the resulting posterior predictive distributions for the total seal pup counts are quite different. This indicates that it is not sufficient to consider the mean behavior. [Salberg et al. \(2009\)](#) found that a negative binomial likelihood was necessary in order to obtain a good fit to the data. However, the conclusions of our analysis indicate that a more careful and comprehensive assessment of both the model uncertainty and the spatial inhomogeneities may warrant the use of the simpler Poisson likelihood, in particular if the doubly stochastic framework of the LGCP is applied.

A direct comparison of the models based on a non-homogeneous versus a doubly stochastic point process formulation is somewhat complicated by the fact that both the models and the estimation approaches differ, and it is thus not possible to separate the effect of the model from the effect of the inference method. In particular, the SPDE-INLA approach is an approximate fully Bayesian approach with priors on all parameters, while the GAM approach uses MPLE and a Bayesian assessment of uncertainty. In particular, we do not include smoothing parameter uncertainty in the GAM approach as we aimed to replicate the implementation of [Salberg et al. \(2009\)](#) for the comparison. While the general GAM methodology has been extended to allow for fully Bayesian inference, see e.g. [Wood \(2016\)](#), such an extension has not yet been implemented in the R package `mgcv` for the negative binomial likelihood. Alternatively, [Umlauf et al. \(2018\)](#) propose a full Bayesian analysis of GAMs with associated software available in the recent R package `bamlss`.

Notably, the predictive distributions for the total seal pup counts are very similar under the homogeneous Poisson model estimated with a fully Bayesian approach and under the non-homogeneous Poisson model estimated with a MPLE/Bayesian uncertainty approach. This invites the conclusion that the vast difference between these two predictive distributions and that under the doubly stochastic Poisson model is, to a large degree, due to the inclusion of the random effect in the intensity function of the doubly stochastic model rather than choice of inference approach.

For the LGCP inference, we have applied the SPDE-INLA approach directly as implemented in the R package `INLA` ([Lindgren et al., 2011](#); [Rue et al., 2009](#)). Alternatively, the more recent package `inlabru` ([Bachl et al., 2019](#)) is based on the SPDE-INLA software to model point pattern data from surveys with varying detection probabilities over the sampled area. Our setting is slightly different, in that the detection probability is assumed to be equal to one over the entire sampled area and we only have sampled counts per each photo rather than the precise locations of the seal pups. While this somewhat simpler setting can also be analyzed with `inlabru`, we have chosen to implement our own version which allows for a slightly higher flexibility in the generation of the mesh. Besides small differences in the manner in which edge effects are treated

when approximating the integral over the latent field, our implementation should give comparable results to `inlabru`.

In our model specification, we only consider linear effects of the satellite covariate and linear spatial effects. We investigated including non-linear effects and squared terms both for the LGCP and the deterministic intensity model formulations. However, as this did not improve the performance of the models, such terms were not included in the final model specifications.

In the present work, harp and hooded seals have been modeled separately, based on a single survey. As an alternative, one may consider building a joint model for harp and hooded seals, for instance by using a common spatial field, in addition to seal specific ones (see e.g. Waagepetersen et al. (2016)). Further, due to drifting ice and moving seals, the spatial locations of the seal pups cannot be directly compared from one survey to the next. However, most of the seals tend to stay in more or less the same packs from one year to another. Such information could potentially be utilized to construct informative priors, which may reduce the modeling uncertainty. It would be interesting to see investigations on such attempts at borrow strength, either from previous surveys or between seal types. Such investigations are, however, out of scope of the present paper.

The analysis in this paper has been carried out in the R programming language (R Core Team, 2019). Code and data associated with the paper are available through the following Github repository: github.com/PointProcess/SealPupProduction-JRSSC-code.

Acknowledgments

Martin Jullum and Thordis Thorarinsdottir acknowledge the support of The Research Council of Norway through grant 240838 “Model selection and model verification for point processes”. Fabian Bachl was supported by EPSRC grant EP/K041053/1 “Modelling spatial distribution and change from wildlife survey data”. We thank Tor Arne Øigård for helpful discussions and for assisting with the photographic survey data which were provided by the Norwegian Institute of Marine Research, Øivind Due Trier for retrieving the satellite image data, Dawid Lawrence Miller and Alex Lenkoski for helpful discussions regarding the various inference approaches.

References

- Bachl, F. E., Lindgren, F., Borchers, D. L. and Illian, J. B. (2019) `inlabru`: an r package for bayesian spatial modelling from ecological survey data. *Methods in Ecology and Evolution*, **10**, 760–766.
- Bellier, E., Monestiez, P. and Guinet, C. (2010) Geostatistical modelling of wildlife populations: a non-stationary hierarchical model for count data. In *Geoenv Vii-geostatistics for environmental applications*, 1–12. Springer.
- Blangiardo, M. and Cameletti, M. (2015) *Spatial and Spatio-temporal Bayesian Models with R - INLA*. Wiley.

- Christensen, O. F. and Waagepetersen, R. (2002) Bayesian prediction of spatial count data using generalized linear mixed models. *Biometrics*, **58**, 280–286.
- Cox, D. R. (1955) Some statistical methods connected with series of events. *Journal of the Royal Statistical Society. Series B (Methodological)*, 129–164.
- De Oliveira, V. (2014) Poisson kriging: A closer investigation. *Spatial Statistics*, **7**, 1–20.
- Diggle, P. J. (2013) *Statistical analysis of spatial and spatio-temporal point patterns*. Chapman and Hall/CRC.
- Diggle, P. J. and Milne, R. K. (1983) Negative binomial quadrat counts and point processes. *Scandinavian journal of statistics*, **10**, 257–267.
- Diggle, P. J., Tawn, J. A. and Moyeed, R. (1998) Model-based geostatistics. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **47**, 299–350.
- Gneiting, T. and Raftery, A. E. (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, **102**, 359–378.
- Good, I. J. (1952) Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 107–114.
- Hastie, T. and Tibshirani, R. (1990) *Generalized additive models*. Wiley Online Library.
- ICES (2014) Report of the ICES/NAFO working group on harp and hooded seals (WGHARP). *Tech. rep.*, ICES CM 2014/ACOM:20, Quebec City, Quebec, Canada.
- Illian, J., Penttinen, A., Stoyan, H. and Stoyan, D. (2008) *Statistical analysis and modelling of spatial point patterns*, vol. 70. John Wiley & Sons.
- 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**, 1189–1210.
- Kovacs, K. M. and Lavigne, D. (1986) *Cystophora cristata*. *Mammalian Species*, 1–9.
- 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.
- Matérn, B. (1971) Doubly stochastic poisson processes in the plane. *Statistical ecology*, **1**, 195–213.
- Matheson, J. E. and Winkler, R. L. (1976) Scoring rules for continuous probability distributions. *Management science*, **22**, 1087–1096.
- Møller, J., Syversveen, A. R. and Waagepetersen, R. P. (1998) Log Gaussian Cox processes. *Scandinavian journal of statistics*, **25**, 451–482.
- Møller, J. and Waagepetersen, R. P. (2003) *Statistical inference and simulation for spatial point processes*. CRC Press.

- 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.
- (2014b) From pup production to quotas: current status of harp seals in the Greenland Sea. *ICES Journal of Marine Science: Journal du Conseil*, **71**, 537–545.
- 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.
- R Core Team (2019) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- 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.
- 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*.
- 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**, 867.
- 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**, 847–858.
- Sergeant, D. E. (1974) A rediscovered whelping population of hooded seals *cystophora cristata* erleben and its possible relationship to other populations. *Polarforschung*, **44**, 1–7.
- (1991) Harp seals, man and ice. *Canadian Journal of Fisheries and Aquatic Sciences*, **114**, 1–153.
- Silverman, B. W. (1985) Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society: Series B (Methodological)*, **47**, 1–21.
- 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**, 49–70.
- 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**, 65–74.

- Simpson, D., Rue, H., Riebler, A., Martins, T. G., Sørbye, S. H. et al. (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, **32**, 1–28.
- Thorarinsdottir, T. L. and Schuhen, N. (2018) Verification: assessment of calibration and accuracy. In *Statistical Postprocessing of Ensemble Forecasts*, 155–186. Elsevier.
- Umlauf, N., Klein, N. and Zeileis, A. (2018) Bamlss: bayesian additive models for location, scale, and shape (and beyond). *Journal of Computational and Graphical Statistics*, **27**, 612–627.
- Waagepetersen, R., Guan, Y., Jalilian, A. and Mateu, J. (2016) Analysis of multispecies point patterns by using multivariate log-gaussian cox processes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **65**, 77–96.
- Wahba, G. (1990) *Spline models for observational data*, vol. 59. Siam.
- Wolpert, R. L. and Ickstadt, K. (1998) Poisson/Gamma random field models for spatial statistics. *Biometrika*, **85**, 251–267.
- Wood, S. (2017) *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2 edn.
- Wood, S. N. (2003) Thin-plate regression splines. *Journal of the Royal Statistical Society (B)*, **65**, 95–114.
- (2006) On confidence intervals for generalized additive models based on penalized regression splines. *Australian & New Zealand Journal of Statistics*, **48**, 445–464.
- (2016) Just another Gibbs additive modeler: Interfacing JAGS and mgcv. *Journal of Statistical Software*, **75**, 1–15.
- Yuan, Y., Bachl, F. E., Lindgren, F., Borchers, D. L., Illian, J. B., Buckland, S. T., Rue, H., Gerrodette, T. et al. (2017) Point process models for spatio-temporal distance sampling data from a large-scale survey of blue whales. *The Annals of Applied Statistics*, **11**, 2270–2297.