

Using spatiotemporal statistical models to estimate animal abundance and infer ecological dynamics from survey counts

PAUL B. CONN,^{1,5} DEVIN S. JOHNSON,¹ JAY M. VER HOEF,¹ MEVIN B. HOOTEN,^{2,3,4} JOSHUA M. LONDON,¹
AND PETER L. BOVENG¹

¹National Marine Mammal Laboratory, NOAA, National Marine Fisheries Service, Alaska Fisheries Science Center,
7600 Sand Point Way Northeast, Seattle, Washington 98115 USA

²U.S. Geological Survey, Colorado Cooperative Fish and Wildlife Research Unit, Colorado State University, Fort Collins,
Colorado 80523 USA

³Department of Fish, Wildlife, and Conservation Biology, Colorado State University, Fort Collins, Colorado 80523 USA

⁴Department of Statistics, Colorado State University, Fort Collins, Colorado 80523 USA

Abstract. Ecologists often fit models to survey data to estimate and explain variation in animal abundance. Such models typically require that animal density remains constant across the landscape where sampling is being conducted, a potentially problematic assumption for animals inhabiting dynamic landscapes or otherwise exhibiting considerable spatiotemporal variation in density. We review several concepts from the burgeoning literature on spatiotemporal statistical models, including the nature of the temporal structure (i.e., descriptive or dynamical) and strategies for dimension reduction to promote computational tractability. We also review several features as they specifically relate to abundance estimation, including boundary conditions, population closure, choice of link function, and extrapolation of predicted relationships to unsampled areas. We then compare a suite of novel and existing spatiotemporal hierarchical models for animal count data that permit animal density to vary over space and time, including formulations motivated by resource selection and allowing for closed populations. We gauge the relative performance (bias, precision, computational demands) of alternative spatiotemporal models when confronted with simulated and real data sets from dynamic animal populations. For the latter, we analyze spotted seal (*Phoca largha*) counts from an aerial survey of the Bering Sea where the quantity and quality of suitable habitat (sea ice) changed dramatically while surveys were being conducted. Simulation analyses suggested that multiple types of spatiotemporal models provide reasonable inference (low positive bias, high precision) about animal abundance, but have potential for overestimating precision. Analysis of spotted seal data indicated that several model formulations, including those based on a log-Gaussian Cox process, had a tendency to overestimate abundance. By contrast, a model that included a population closure assumption and a scale prior on total abundance produced estimates that largely conformed to our a priori expectation. Although care must be taken to tailor models to match the study population and survey data available, we argue that hierarchical spatiotemporal statistical models represent a powerful way forward for estimating abundance and explaining variation in the distribution of dynamical populations.

Key words: abundance; Bering Sea; count data; dimension reduction; *Phoca largha*; process convolution; resource selection; spatiotemporal model; spotted seal.

INTRODUCTION

A tacit assumption in many models for animal count data is that animal density remains constant over the landscape while sampling is being conducted. This assumption is clearly unrealistic for populations that inhabit dynamic landscapes or whose spatial distribution varies considerably over the course of a study. For instance, the distribution and abundance of bird species inhabiting agricultural landscapes can vary substantially within seasons due to cultivation and harvesting practices (Miller et al. 2011), as can abundance of bird

and amphibian populations that are dependent on dynamic wetland environments (Murkin et al. 1997, Babbitt and Tanner 2000). Our own interest in spatiotemporal variation in animal abundance stems from involvement in surveys of ice-associated seals, where the quantity and distribution of suitable habitat (seasonal sea ice) changes dramatically on the scale of weeks or even days (Ver Hoef and Jansen 2007, Ver Hoef et al. 2014).

Early development of models for analysis of animal count data (see Buckland et al. [2001], and references therein) focused on design-based estimation (cf. Cochran 1977), which is perhaps the most simple and reliable approach for obtaining snapshot estimates of abun-

Manuscript received 21 May 2014; revised 29 August 2014;
accepted 14 October 2014. Corresponding Editor: D. C. Speirs.

⁵ E-mail: paul.conn@noaa.gov

dance, provided that the spatial configuration of animal density does not vary appreciably while the study is being conducted and that a sampling frame (e.g., systematic, stratified random samples) can be developed and strictly adhered to. More recently, model-based approaches for analyzing animal abundance from transect or plot-based counts (Hedley and Buckland 2004, Ver Hoef and Jansen 2007, Johnson et al. 2010) have become more prevalent. Although these approaches require additional (often distributional) assumptions, model-based inference permits greater flexibility because sampling effort can be allocated at different times and locations than originally planned, and because features such as temporal (Moore and Barlow 2011) and spatial (Miller et al. 2013) autocorrelation can be included and visualized. Model-based formulations are also much better suited for allowing animal density to be expressed as a function of habitat covariates, allowing ecologists to make inferences about factors that explain variation in animal abundance. Such inferences are crucial to both the conservation of wild populations and the study of their ecology, as exemplified by the recent surge in species distribution modeling research (SDMs; Guisan and Thuiller 2005, Elith and Leathwick 2009). In this study, we focus exclusively on count data, although we note that many of the same modeling constructs can also be applied to presence–absence data (as commonly used in SDMs).

Development of model-based estimation procedures for analysis of count data have largely focused on inclusion of spatial variation in abundance without as much consideration of temporal variation. Notable exceptions include Ver Hoef and Jansen (2007), who included changing habitat covariates and additive spatial and temporal effects in count models for harbor seals; Fink et al. (2010), who developed machine learning algorithms to explore spatiotemporal variation in large-scale migratory bird survey data; Miller et al. (2011), who allowed bird density to rely on habitat covariates that varied between sample occasions (using an N -mixture formulation; see Royle et al. 2004); Moore and Barlow (2011), who allowed the modeled density of fin whales to vary among years and spatial strata; and Ver Hoef et al. (2014), who conceptualized phocid seal counts as a function of both changing habitat covariates and spatially autocorrelated random effects. Hierarchical generalized linear mixed models incorporating both spatial and temporal effects are also prevalent in analysis of population trends from count data (Sauer and Link 2011, Ross et al. 2012), but it is not often possible to scale up from these models to the estimates of absolute density or abundance that are frequently needed for population management.

The statistical literature on spatiotemporal models has exploded in recent years (see Cressie and Wikle [2011] for a review), and in theory, it would be possible to embed a large number of spatiotemporal models for animal density into model-based estimation procedures

for transect or plot-based counts. Cressie and Wikle (2011) categorize such models as either descriptive, where spatiotemporal autocorrelation is codified by allowing positive covariance at times and locations that are close together, or dynamical (i.e., process-based), where the investigator uses knowledge of the study system to help specify a dynamical model governing the evolution of the spatiotemporal process over time (e.g., via difference equations; Wikle and Hooten 2010). Spatiotemporal models often include one or more strategies for dimension reduction to ease computational burden (of which there are many; for a review see Wikle [2010]). As such, it will not be possible to enumerate all possible approaches for incorporating spatiotemporal structure into abundance models. Instead, we start by reviewing recent work in this field, aiming to provide ecologists with intuition on the flavors of spatiotemporal models currently available in the literature, providing connections with relevant ecological theory where applicable (e.g., with ideal free distributions; Fretwell and Lucas 1970).

Next, we turn our focus to spatiotemporal modeling considerations as they relate specifically to abundance estimation. These include boundary conditions (emigration/immigration from the study area), demographic closure (population increases/decreases due to mortality or reproduction), choice of link function, and extrapolation of fitted relationships to unsampled areas. After delineating a (by no means exhaustive) suite of spatiotemporal models for abundance estimation, including novel formulations embodying resource selection and permitting a form of population closure, we illustrate the consequences of different modeling choices by analyzing both simulated data and seal count data from transects over the Bering Sea. In both cases, the quantity and quality of available habitat changes dramatically while surveys are being conducted.

SPATIOTEMPORAL MODELS FOR COUNT DATA

Animal count data are most often obtained on areal sample units (i.e., in discretized space) and at discrete points in time. Even if sampling is not instantaneous, a common convention in spatiotemporal models is to discretize time at some level (e.g., daily) and to treat all counts gathered in a given day as if they occurred instantaneously. We follow these conventions, and model both the abundance of animals and observed counts by appealing to specific points in time $t \in \{1, 2, \dots, T\}$ and specific sample units $s \in \{1, 2, \dots, S\}$. For simplicity, we assume that all sample units are the same size (this requires apportioning counts into respective sample units whenever transect sampling is conducted). Our task is to construct plausible spatiotemporal models that describe how a sample of observed counts at a subset of n locations and times $\{(s_1, t_1), (s_2, t_2), \dots, (s_n, t_n)\}$ are related to absolute abundance $N_{s,t}$ at all times and sample units.

All spatiotemporal models that we examine relate observed counts, $C_{s,t}$, obtained at a limited number of times and sample units, to a hidden spatiotemporal process $\mu_{s,t}$ that is defined over all sample units and times through a standard probability mass function (PMF), link function, covariates \mathbf{x} , and parameters $\boldsymbol{\theta}$. For example, the formulation

$$\begin{aligned} [\mathbf{C}] &= \text{Poisson}(\boldsymbol{\lambda}) \\ \boldsymbol{\lambda} &= \exp(\mathbf{o} + \mathbf{H}\boldsymbol{\mu}) \end{aligned} \quad (1)$$

specifies that counts are Poisson distributed with a log-link function on latent abundance intensity ($\boldsymbol{\lambda}$). Following Wikle and Hooten (2010), we use the matrix \mathbf{H} to map from the full set of ST values of $\mu_{s,t}$ to the smaller set of n counts. We also include log offsets \mathbf{o} to allow for differences in the amount of area surveyed (this could also be used to help model detection probability when detectability is less than 1.0; see *Discussion*). In practice, more flexible distributional forms could be considered, such as the negative binomial or Conway-Maxwell Poisson (Wu et al. 2013), but our focus here is in describing general forms for the spatiotemporal process $\mu_{s,t}$, and we often use the Poisson as a convenient PMF.

Models with the form of Eq. 1 can sometimes be analyzed directly via maximum likelihood, but introducing even moderate complexity in models for $\mu_{s,t}$ can make computation intractable. To preserve flexibility, we prefer to specify prior distributions for model parameters and to conduct Bayesian inference (Gelman et al. 2004) on complete data likelihoods (Dempster et al. 1977). We follow the tradition of Berliner (1996) in expressing the posterior distribution associated with Eq. 1 as consisting of an observation model, a process model, and a set of prior distributions. Letting $[a|b]$ denote the conditional distribution of a given b , we write the joint posterior distribution for parameters $\boldsymbol{\theta}$ and latent variables $\boldsymbol{\mu}$ (up to proportionality constant) as

$$[\boldsymbol{\mu}, \boldsymbol{\theta} | \mathbf{C}, \mathbf{o}] \propto [\mathbf{C} | \mathbf{o}, \boldsymbol{\mu}] [\boldsymbol{\mu} | \boldsymbol{\theta}] [\boldsymbol{\theta}]. \quad (2)$$

Here, $[\mathbf{C} | \mathbf{o}, \boldsymbol{\mu}]$ simply describes the Poisson observation model (Eq. 1), $[\boldsymbol{\mu} | \boldsymbol{\theta}]$ describes the process model (how latent abundance varies over space and time), and $[\boldsymbol{\theta}]$ describes the prior distribution on model parameters. Although care must be taken in formulating observation and prior parameter models, we now turn to describing different ways of formulating the spatiotemporal process model before discussing ways of using dimension reduction to decrease computational burden.

Descriptive vs. dynamical formulations

Cressie and Wikle (2011) classify spatiotemporal models into two types: descriptive and dynamical. The development of descriptive spatiotemporal models largely follows the development of spatial and time-series models, where there is an a priori expectation that system states (in our case, $\mu_{s,t}$) close together in space or

time should be similar to one another. In descriptive models, this notion is often codified by specifying a covariance function that decreases as a function of distance (where distance can be either in space or time). For instance, the model

$$\mu_{s,t} = \mathbf{X}_t \boldsymbol{\beta} + \eta_s + \gamma_t + \kappa_{s,t} + \varepsilon_{s,t} \quad (3)$$

might be used to indicate that the expected number of animals at location s at time t is a function of a design matrix \mathbf{X}_t at time t (consisting of covariate values, etc.) and regression parameters $\boldsymbol{\beta}$, a purely spatial random effect (η_s), a purely temporal random effect (γ_t), a spatiotemporal random effect ($\kappa_{s,t}$), and independent random error $\varepsilon_{s,t}$. Statisticians have developed a large number of approaches to model autocorrelation in random effects as purely spatial (e.g., geostatistical models, Gaussian Markov random fields [GMRFs]; Cressie 1993, Banerjee et al. 2004, Rue and Held 2005) or temporal (e.g., AR[p] processes and ARIMA models; Box et al. 2008, Prado and West 2010). Similar approaches can be used to define spatiotemporal covariance functions (Cressie and Wikle 2011) for the $\kappa_{s,t}$ random effects.

By contrast, dynamical spatiotemporal models attempt to model the mechanism by which the spatiotemporal process evolves over time. In spatiotemporal models for physical processes (e.g., atmospheric processes), the evolution of the spatiotemporal process is largely governed by the laws of physics, which can be represented with ordinary or partial differential equations (Wikle et al. 2001, Wikle and Hooten 2010). Discrete time approximations to such equations can provide much more meaningful structure than simple descriptive models, although descriptive model features can still be used to account for residual spatiotemporal autocorrelation. In ecology, dynamical spatiotemporal models have been used to model invasive species spread (Wikle 2003, Hooten et al. 2007, Hooten and Wikle 2008). Even though invasive species do not necessarily subscribe to a single succinct law that governs their spread, physical models for advection–diffusion processes are still very useful for describing and constraining evolution of the spatiotemporal process.

Returning to the problem of estimating animal abundance, consider a dynamical model for $\mu_{s,t}$ that takes on the form

$$[\boldsymbol{\mu} | \boldsymbol{\theta}] = [\boldsymbol{\mu}_1 | \boldsymbol{\theta}] \prod_{t=2}^T [\boldsymbol{\mu}_t | \boldsymbol{\mu}_{t-1}, \boldsymbol{\theta}] \quad (4)$$

where subscripts in this case denote time (so that $\boldsymbol{\mu}_t$ describes a collection of $\mu_{s,t}$ values over all spatial locations for a given time t). With this setup, we require a model for the initial spatial distribution of $\mu_{s,1}$ at the beginning of the study (i.e., $[\boldsymbol{\mu}_1 | \boldsymbol{\theta}]$), and specify a dynamical, Markovian model for evolution of the process via $[\boldsymbol{\mu}_t | \boldsymbol{\mu}_{t-1}, \boldsymbol{\theta}]$. For instance, we might consider a matrix model, e.g.,

$$\boldsymbol{\mu}_t = \mathbf{M}_t \boldsymbol{\mu}_{t-1} + \boldsymbol{\varepsilon}_t \quad (5)$$

where the matrix \mathbf{M}_t is possibly a function of covariates at time t , \mathbf{x}_t , and (unknown) parameter values $\boldsymbol{\theta}$ that describe how animals redistribute themselves over the landscape from $t \rightarrow t+1$. A variety of processes could be used to aid construction of \mathbf{M}_t , including advection–diffusion (Wikle 2003) and agent-based models (Hooten and Wikle 2010). We explore a hybrid of these approaches that incorporates resource selection via redistribution kernels in *Spatiotemporal abundance models*.

In an ecological context, descriptive models applied to abundance estimation essentially reflect ideal free distributions (Fretwell and Lucas 1970), whereby animal abundance is free to shift to preferable habitat between time steps, regardless of how far apart such habitat is from animal abundance epicenters in previous time steps. Although inclusion of spatiotemporal random effects have potential to prevent large shifts in abundance from time step to time step, such artifacts may be expected in descriptive statistical models, especially when survey data are sparse. By contrast, dynamical models have the capacity to constrain shifts in abundance from time step to time step.

Dimension reduction

When spatial models include thousands of locations, the computational demands involved in fitting full-dimensional models can be prohibitive. Estimation of spatial and spatiotemporal models inevitably requires evaluating the inverse of covariance matrices for random effects, the calculation of which requires operations that increase as a cubic function of the number of locations (Wikle 2010). Although specification of sparse covariance matrices or GMRF models can help reduce the number of calculations required, it is not always sufficient to obtain a solution in a reasonable amount of time. In particular, Bayesian analysis of hierarchical models via Markov chain Monte Carlo (MCMC) often requires hundreds of thousands of matrix inversions. Inclusion of a temporal dimension on top of spatial structure only serves to increase dimensionality, further increasing computational requirements.

A common approach to this predicament in high-dimension spatial and spatiotemporal models is to employ some sort of dimension reduction strategy. Dimension reduction includes a reduction in the number of effective random effect parameters as well as the size of associated covariance matrices that need to be inverted during estimation. There are numerous strategies for dimension reduction, but in each we express a high-dimensional set of random effects (say $\boldsymbol{\eta}$ from Eq. 3) as the product of an expansion matrix, \mathbf{K} , and a reduced parameter set $\boldsymbol{\alpha}$. For instance, we may write

$$\boldsymbol{\eta} = \mathbf{K}\boldsymbol{\alpha} \quad (6)$$

where the $m \times 1$ dimensional vector $\boldsymbol{\alpha}$ holds spatial random effects on a reduced parameter space, and the $S \times m$ dimensional matrix \mathbf{K} maps these reduced dimension random effects to S -dimensional space. The consequence of such a formulation is a reduction in parameter space from $S \rightarrow m$, which can greatly improve computational tractability.

There are many approaches for parameterizing \mathbf{K} ; for instance, the rows of \mathbf{K} can be specified using any of a number of orthogonal or kernel basis functions (Wikle 2010), and can potentially include additional parameters (e.g., kernel bandwidth). Similarly, different possibilities exist for $\boldsymbol{\alpha}$. In certain knot-based approaches to dimension reduction such as predictive process models (Banerjee et al. 2008), the $\boldsymbol{\alpha}$ are spatially referenced, while in others (e.g., spectral decomposition) they are not.

Throughout this study, we focus on a particular form of dimension reduction that uses kernel convolution (also termed process convolution; Barry and Ver Hoef 1996, Higdon 1998) to interpolate spatiotemporal random effects relative to m knots that are distributed across space and/or time. This allows us to model spatiotemporal random effects at the knot level, and since $m \ll S$, to markedly improve computational efficiency. In this case \mathbf{K} is simply used to interpolate across knots, with entries defined as

$$K_{s,m} = w_{s,m} / \sum_s w_{s,m} \quad (7)$$

where

$$w_{s,m} = (2\pi\sigma^2)^{-1} \exp\left(\frac{-d_{s,m}^2}{2\sigma^2}\right) \quad (8)$$

and $d_{s,m}$ gives the Euclidean distance between the centroid of sample unit s and the location of knot m (where Euclidean distance is calculated relative to a geographical projection of the Earth's surface onto a two-dimensional map). Alternate kernels are certainly possible, and would induce alternate covariance structure. Theoretically, it is possible to estimate optimal kernel bandwidths (σ), and to consider refinements allowing for anisotropy (Higdon 1998). However, for the present development, we limit ourselves to the simple isotropic specification in Eq. 8, and use a rule of thumb that sets σ equal to the distance between nearest neighbor knots (Higdon 2002). This approach has been successfully employed for dimension reduction in applied spatiotemporal settings, such as modeling ozone concentrations over space and time (Calder et al. 2002, Higdon 2002). We use two different specifications for $\boldsymbol{\alpha}$ in subsequent models, depending on whether we are working with spatial or spatiotemporal random effects.

Abundance estimation considerations

Estimation of animal abundance introduces a number of nuances and important considerations above and beyond those associated with simply fitting a statistical model. These considerations differ depending on whether one selects a purely descriptive or a purely dynamic model for the evolution of the spatial process. Hybrid models thus require making important decisions about all of the following points.

Boundary conditions.—When a mechanistic (dynamic) model is used to help guide the evolution of the spatiotemporal process, one must pay attention to the boundary conditions necessary to fully specify the solution to systems of partial differential equations (Haberman 1998). With respect to animal abundance, boundary conditions specify how animals immigrate and emigrate across the edges of the study area. For instance, the dynamic matrix model formulation in Eq. 5 could be expanded to include separate consideration of movement within the study area and movement across boundaries

$$[\boldsymbol{\mu}_t | \boldsymbol{\mu}_{t-1}, \boldsymbol{\theta}] = \mathbf{M}_t \boldsymbol{\mu}_{t-1} + \mathbf{M}_t^b \boldsymbol{\mu}_{t-1}^b + \boldsymbol{\varepsilon}_t$$

where the superscript b denotes cells along the boundary of the study area (Cressie and Wikle 2011, Section 6.3.2). It is possible to model discretized fluxes across boundaries in hierarchical models (Wikle et al. 2003), and such extensions will be necessary in certain applications (for instance, when estimating abundance of migrating animals). However, for the present application, we make the simplification that immigration and emigration from the study area are zero. This is both for clarity of presentation, and because a previous analysis (Ver Hoef et al. 2014) of our focal seal population suggested that total abundance was relatively constant in the study area while surveys were being conducted.

Demographic closure and choice of link function.—The rate at which the target population increases or decreases due to mortality and births may help guide selection of an appropriate spatiotemporal statistical model. For instance, invasive species often spread by both range extension and increasing abundance in previously occupied sites. As such, related spatiotemporal models have often included a Malthusian growth parameter (Wikle 2003, Hooten and Wikle 2008) to allow abundance to increase over the course of the study. On the other hand, if the population is stable while the study is being conducted, use of a population closure assumption (a conservation of mass property) may help to stabilize estimation and to provide a single estimate of abundance (as opposed to a sequence of time-referenced estimates). This notion is perhaps easier stated than implemented, however. First, it is difficult to implement such a restriction in purely descriptive statistical models (e.g., using a log-Gaussian Cox process formulation such as Eq. 1). Second, a log-link

function for $\mu_{s,t}$ can also make this constraint difficult to enforce in dynamic models. For example, it is a simple matter to enforce $\sum_s \mu_{s,t} = \sum_s \mu_{s,t+1} \forall t$, but this constraint alone will lead to expected population increases in the case of range contraction, or expected population decreases in the case of range expansion if $\mu_{s,t}$ is modeled in log space. Dynamic models for spatiotemporal count data have largely been specified in log space (e.g., Wikle 2003, Hooten and Wikle 2008); in these studies, we expect inclusion of a Malthusian growth parameter to help balance out expected population decreases during range expansion. We examine several approaches for inducing a conservation of mass property (see *Spatiotemporal abundance models*), and also investigate whether this seemingly innocuous modeling choice (the form of link function for $\mu_{s,t}$) has ramifications for abundance estimation.

Spatiotemporal extrapolation.—If we were simply interested in making inferences about factors that influence variation of counts in times and places where sampling occurred, the model specified in Eq. 1 might suffice, and be all we would need to describe how the mean abundance intensity $\mu_{s,t}$ varies in time and space. A simple descriptive model for $\mu_{s,t}$ might include habitat effects and spatiotemporal autocorrelation, several forms of which can be implemented in a generalized linear mixed-modeling framework (Ross et al. 2012). However, since we are concerned with estimation of total abundance, we must also make inferences about how $\mu_{s,t}$ varies in times and locations where we do not sample.

The standard procedure for generating abundance estimates from hierarchical models is to use posterior prediction to simulate abundance in unsampled areas, and to combine these with known counts from sampled areas (Ver Hoef 2008, Johnson et al. 2010) to generate an estimate of total abundance. This formulation is attractive in that it includes a finite population correction; if the entire study is sampled and all animals are detectable, abundance is known with certainty. Unfortunately, generating posterior predictions over space and time is not without its perils. In addition to the usual caution of not extrapolating past the range of observed data, inclusion of descriptive spatial structure (e.g., via conditionally autoregressive models) can sometimes lead to spuriously high predictions of abundance in locations that are near edges of the study area (edge effects; Ver Hoef and Jansen 2007) or in locations with large gaps in spatial coverage (Conn et al. 2014). Although alternate posterior estimators can help (e.g., the Varian predictor induced by the linex loss function; Ver Hoef and Jansen 2007), there seems to be a clear need for tailoring the complexity of spatiotemporal models (particularly descriptive models) to match the richness of the data set. This is especially the case when using a log-link function; as prediction error increases, estimates of abundance are likely to increase as well, simply because the error distribution is right-

skewed (Ver Hoef 2008). As such, it is important to conduct simulation exercises and analyze test data sets to help establish the robustness of particular model structures in yielding robust and defensible abundance estimates.

SPATIOTEMPORAL ABUNDANCE MODELS

We consider several different parameterizations for descriptive and dynamical models and examine their performance in analyzing simulated and real data sets.

Additive space–time (AST) formulation

The first approach we describe is a simple descriptive model for $\mu_{s,t}$ that includes additive spatial and temporal structure (i.e., omitting space–time interactions). To induce this structure, we modify Eq. 3 as

$$\mu_{s,t} = \mathbf{X}_{s,t}\boldsymbol{\beta} + \eta_s + \gamma_t + \varepsilon_{s,t}.$$

The implication of this additive construction is that the temporal random effect γ_t will lead to uniform increases or decreases in log abundance over the entire landscape. This formulation may not be realistic for many dynamical animal populations, but may still be useful in cases where resources remain relatively constant, dynamics are governed primarily by birth–death processes, and/or survey data are sparse. A similar formulation was adopted by Ver Hoef and Jansen (2007) to estimate abundance of harbor seals.

We model spatial structure by distributing knots across the landscape and imposing a process convolution model for η_s as in Eq. 6. In this case, the random spatial effects, $\boldsymbol{\alpha}$, represent weights associated with each knot location, and are given independent Gaussian prior distributions. For γ_t , we specify an RW2 prior distribution, which is a second-order Markov random walk time-series process (Rue and Held 2005). This formulation imposes a smoother temporal trend in log of latent abundance than would a first-order Markov random walk (or RW1) process. We describe a Bayesian implementation of this model in Appendix A.

A spatiotemporal process convolution (STPC) model

The next model we implemented shares many similarities with the AST model, but allows for spatiotemporal interactions. Also a descriptive (as opposed to dynamical formulation), we start with the same general structure and observation model as, i.e., given by Eqs. 1 and 2. However, in this case, we write log of abundance intensity as a function of a single spatiotemporal effect, $\kappa_{s,t}$

$$\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\kappa} + \boldsymbol{\varepsilon}.$$

We will write the $\boldsymbol{\kappa}$ as a function of spatiotemporal parameters on a reduced dimensional space associated with m knots placed across the landscape, such that

$$\boldsymbol{\kappa} = \mathbf{L}\boldsymbol{\alpha}.$$

Using this setup, there are mT $\alpha_{k,t}$ random effects (one for each knot and time step), and \mathbf{L} is a sparse ($ST \times mT$) matrix, whose nonzero elements consist of rearranged entries from the \mathbf{K} matrix. We allow the $\alpha_{k,t}$ parameters to change smoothly over time by imposing an RW2 prior on $\{\alpha_{k,1}, \alpha_{k,2}, \dots, \alpha_{k,T}\}$ for each knot k . This is a conceptually similar approach to that used by Calder et al. (2002) and Higdon (2002) to describe space–time structure in ozone concentrations, and is described further in Appendix A.

Resource selection model

We implemented a novel spatiotemporal model incorporating resource selection as a dynamical process. Our rationale was that if the quantity and quality of available habitat change over time, animals should be able to select areas of higher quality, subject to movement constraints. We followed the general structure for dynamical models for $\mu_{s,t}$ presented in Eq. 4 (with a subtle modification to improve computational tractability; see Appendix A) and model latent abundance at the beginning of the study using a descriptive formulation

$$\boldsymbol{\mu}_1 = \mathbf{X}_1\boldsymbol{\beta}_1 + \boldsymbol{\eta} + \boldsymbol{\varepsilon}_1$$

where $\boldsymbol{\eta}$ are modeled using a process convolution formulation as in the AST model and the 1 subscript indicates covariates and parameters specific to the first time step of the study. For future time steps, we let latent abundance redistribute itself according to the weighted distribution commonly used in resource selection applications (cf. Patil and Rao 1978, Lele and Keim 2006, Johnson et al. 2008). In typical resource selection studies employing use–availability designs, ecologists are interested in explaining variation in an animal’s use of spatial locations relative to the set of spatial locations that were available to them. In a discrete-space use–availability setting, use of the weighted distribution implies that the probability of transitioning from sample unit $a \in S$ at time $t - 1$ to sample unit $b \in S$ at t is

$$\psi_t^{ab} = \frac{w_{b,t}\varphi_{a,b}}{\sum_s w_{s,t}\varphi_{a,s}}. \quad (9)$$

This probability consists of two components: the relative suitability of sample unit b relative to other sample units at time t ($w_{s,t}$; often expressed as a function of habitat covariates), and the relative connectivity of locations a and b , $\varphi_{a,b}$.

Practically, the $\varphi_{a,b}$ compose a redistribution kernel in absence of habitat selection, and can be represented with a bivariate distribution. For instance, modeling $\varphi_{a,b}$ using a bivariate Gaussian distribution with diagonal covariance structure is indicative of isotropic, diffusive movement in absence of a gradient in habitat quality. According to this approach (which we use throughout the study):

$$\varphi_{a,b} \propto \text{Normal}(d(a,b), \sigma_d^2) \quad (10)$$

where $d(a, b)$ specifies the Euclidean distance between the centroids of sample units a and b . For all examples in this study, we used the following steps to calculate $\varphi_{a,b}$: (1) Set $\varphi_{a,b} = 0$ for all sample units b that do not intersect with a circle that has a radius of three sample units symmetric about the centroid of sample unit a (see Fig. 1). (2) For remaining cells, calculate $\varphi_{a,b}$ as in Eq. 10, using centroids of cells a and b to calculate $d(a, b)$.

This approach is useful computationally, in that the $(S \times S)$ matrix $\boldsymbol{\Phi}$ is sparse and is made up of just nine different numeric values. Our specification of three sample units was initially developed for anticipated use with spotted seal data; in particular, examination of satellite telemetry records indicated that the probability of spotted seals moving >3 sampling units (75 km) in a given day was quite low. The width of the redistribution kernel will likely need to be revisited in different real-world applications, as it is linked to both the area of relative sample units and the dispersal propensity of the species being modeled.

The other component of the transition probabilities in Eq. 9 is $w_{s,t}$, which specifies the suitability of sample unit s at time t relative to other sample units. In the present development, we specify a log-linear model for $w_{s,t}$, where

$$\log(w_t) = \mathbf{X}_t \boldsymbol{\beta}$$

letting habitat covariates (but not regression parameters) vary over time. In contrast to purely descriptive models for $\mu_{s,t}$, we do not include an intercept in this model, as it cancels out of Eq. 9.

Specifying a model for ψ_t^{ab} gives us a means to apply resource selection ideas to guide evolution of the $\mu_{s,t}$. For instance, abundance at time t in a location s is a function of both (1) latent abundance that was in sample unit s at time $t - 1$ that remained in s , and (2) latent abundance in sample units near s that transitioned into s between t and $t + 1$. However, we have yet to specify a complete model for how such transitions work. In the present application, we set

$$\boldsymbol{\mu}_t = \mathbf{M}_t \boldsymbol{\mu}_{t-1} + \boldsymbol{\gamma}_t + \boldsymbol{\varepsilon}_t$$

as a specification we will refer to as open-population resource selection (OPRS). Note that for the OPRS model, the elements of the transition matrix \mathbf{M}_t are simply composed of the transition probabilities ψ_t^{ab} . As with the AST and STPC models, we still use the specification in Eq. 1 to relate μ_t to observed counts. Since $\mu_{s,t}$ is modeled in log space, the OPRS specification means that transitions are specified relative to the natural log of abundance intensity. We attempt to compensate for the artifactual effects of range contraction and expansion on OPRS modeled abundance by including a time-specific random effect, $\boldsymbol{\gamma}_t$, which we model using an AR(1) formulation; AR(1) is an

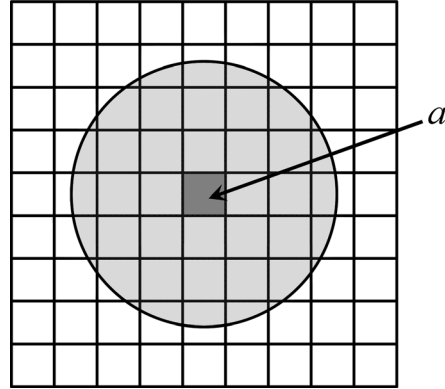


FIG. 1. A visual depiction of the process used to select cells that were “available” in spatiotemporal abundance models incorporating resource selection. We set $\varphi_{a,b}$ values (relative connectivity between sample units a and b) to zero for all sample units b that did not intersect with a circle that has a radius of three sample units symmetric about the centroid of sample unit a (dark-gray shaded cell). Light gray cells indicate sample units for which $\varphi_{a,b} > 0$.

autoregressive model used in time-series analysis. In theory, it would be possible to also consider a closed population analogue of this model where transitions are modeled on the real (as opposed to the log) scale, but it would be difficult to include Gaussian error in this case because abundance intensity must be strictly positive. We leave development of such a model to future investigation.

A closed-population ideal free (CPIF) model

The OPRS implementation provides a possible approach for implementing a pseudo-closure assumption, but may be too limiting when animal dispersal is either leptokurtic or “fast” in relation to the time step of the model. As another approach for constraining expected abundance to be constant over time, we consider a closed-population descriptive (ideal free) model where abundance is reallocated to sample units at each time step proportionally to habitat suitability. In this case, we depart from the general formulation of Eq. 2. Instead, we consider the model

$$[N, \mathbf{N}, \boldsymbol{\omega}, \boldsymbol{\theta} | \mathbf{C}, \mathbf{o}] \propto [\mathbf{C} | \mathbf{N}, \mathbf{o}, \boldsymbol{\omega}] [\boldsymbol{\omega} | \boldsymbol{\theta}, \mathbf{o}] [\mathbf{N} | N, \boldsymbol{\omega}] [\boldsymbol{\theta}]. \quad (11)$$

In particular, we model the counts \mathbf{C} as arising from a binomial sampling process:

$$C_{s,t} \sim \text{binomial}(N_{s,t}, p_{st})$$

where p_{st} is the fraction of the sample unit s that is surveyed at time t and is assumed known (see *Discussion* for possible adaptations to modify p_{st} for incomplete detection). We model abundance in each sample unit s at time t , $N_{s,t}$ as

$$\{N_{1,t}, N_{2,t}, \dots, N_{s,t}\} \sim \text{multinomial}(N, \boldsymbol{\pi})$$

where

$$\pi_{i,t} = \frac{\exp(\omega_{i,t})}{\sum_s \exp(\omega_{s,t})}$$

We express habitat preference ($\omega_{s,t}$) values using a log-linear formulation that includes spatiotemporal autocorrelation

$$\omega_t = \mathbf{o} + \mathbf{X}_t \boldsymbol{\beta} + \boldsymbol{\kappa}_t + \boldsymbol{\varepsilon}_t.$$

Here, \mathbf{o} denotes offsets accounting for possible differences in the size of sample units; by convention, we do not include an intercept in the regression part of the model, as it is not identifiable. As in the spatiotemporal process convolution model, we include temporally evolving spatiotemporal random effects $\boldsymbol{\kappa}_t$ by convolving bivariate normal kernels where the kernel weights change according to an RW2 process (see Appendix A for further details).

COMPUTING

We developed custom-built MCMC samplers in the R programming environment (R Development Core Team 2013) to compute posterior summaries for the models described in *Spatiotemporal abundance models*. Mathematical details of these samplers, including full conditional distributions and prior distribution specifications, are provided in Appendix A. We developed an R package, “STabundance,” which includes all functions, scripts, and data sets to recreate simulation and spotted seal example analyses in this paper. This package is has been archived and published (Conn 2014; active releases *available online*).⁶

GENERIC SIMULATION STUDY

Simulation design

In order to examine the utility of each of our proposed modeling approaches, we conducted a small factorial simulation study, which varied by the type of abundance-generating model (three levels) and the amount of sampling effort (two levels). Each simulation involved sampling and analysis of animal abundance over a 20×20 grid for 20 time steps (e.g., Fig. 2), and was conducted using the following sequence of steps: (1) simulate spatiotemporal variation in a hypothetical habitat covariate across a landscape, (2) simulate spatiotemporal animal abundance across the simulated landscape according to the specified data-generating structure, (3) simulate transect placement and transect counts across each landscape and time step, and (4) conduct Bayesian inference with each of our estimation models (AST, STPC, OPRS, CPIF).

We now briefly describe each of the enumerated steps in turn. A fuller description is provided in Appendix B.

Generating a spatiotemporal habitat covariate.—For each simulation, we used a 30×30 grid for configuring habitat composition. We induced spatiotemporal varia-

tion in a single habitat covariate via dynamic process convolution (Calder et al. 2002, e.g., Fig. 2).

Abundance-generating models.—We used three different models to generate spatiotemporal count data from simulated transect surveys. These included (1) a closed-population model with restricted dispersal (CPRD), (2) a closed-population model with unlimited dispersal (CPUD), using the previously defined CPIF structure, and (3) an open population with restricted dispersal (OPRD).

The OPRD model used the CPRD structure to generate data, but on a larger grid (i.e., 30×30) than was used for analysis (i.e., 20×20). By contrast, CPRD and CPUD used a smaller, 20×20 grid (centered in the middle of the 30×30 grid) to simulate abundance. In effect, the OPRD-generating model allows abundance to shift in and out of the area being surveyed, while CPRD and CPUD do not. For each generating model, simulations were configured such that animals had preference for intermediate values of the habitat covariate (see Appendix B).

Simulating transect surveys and animal counts.—At each time step of simulation, transects were simulated as straight, vertical lines across the entire study area, beginning at randomly selected grid cells along the southern terminus (see e.g., Fig. 2). Each transect thus covered 20 grid cells, and we configured simulations such that transects covered 5% of each grid cell they passed through. Transect placement was re-randomized at each time step. Following selection of the grid cells and times where sampling occurred, we drew animal counts from a Poisson distribution at times and locations where sampling was conducted

$$C_{s,t} \sim \text{Poisson}(0.05\lambda_{s,t}).$$

Estimating abundance from simulated data.—Depending upon simulation configuration, we used anywhere between one and four estimation models (see Appendix B) to estimate animal abundance for each data set. These included AST, CPIF, OPRS, and STPC (for further detail on these models and prior distributions used, see Appendix A). Each model was provided with the spatiotemporal covariates used in abundance generation for use in the linear predictor of each estimation model. Following completion of all simulation runs, we tabulated proportion relative bias, coefficient of variation (CV), and mean squared error of abundance estimates. Calculations were performed slightly differently depending on whether simulation and estimation models were open or closed. For instance, for three of the estimation models (AST, OPRS, and STPC), there were separate posterior distributions of abundance for each time step. To produce a single, “best” average estimate of abundance (\bar{N}) over the whole time frame, we used numerical optimization to solve for the \bar{N} that minimizes

⁶ <https://github.com/pconn/STabundance/releases>

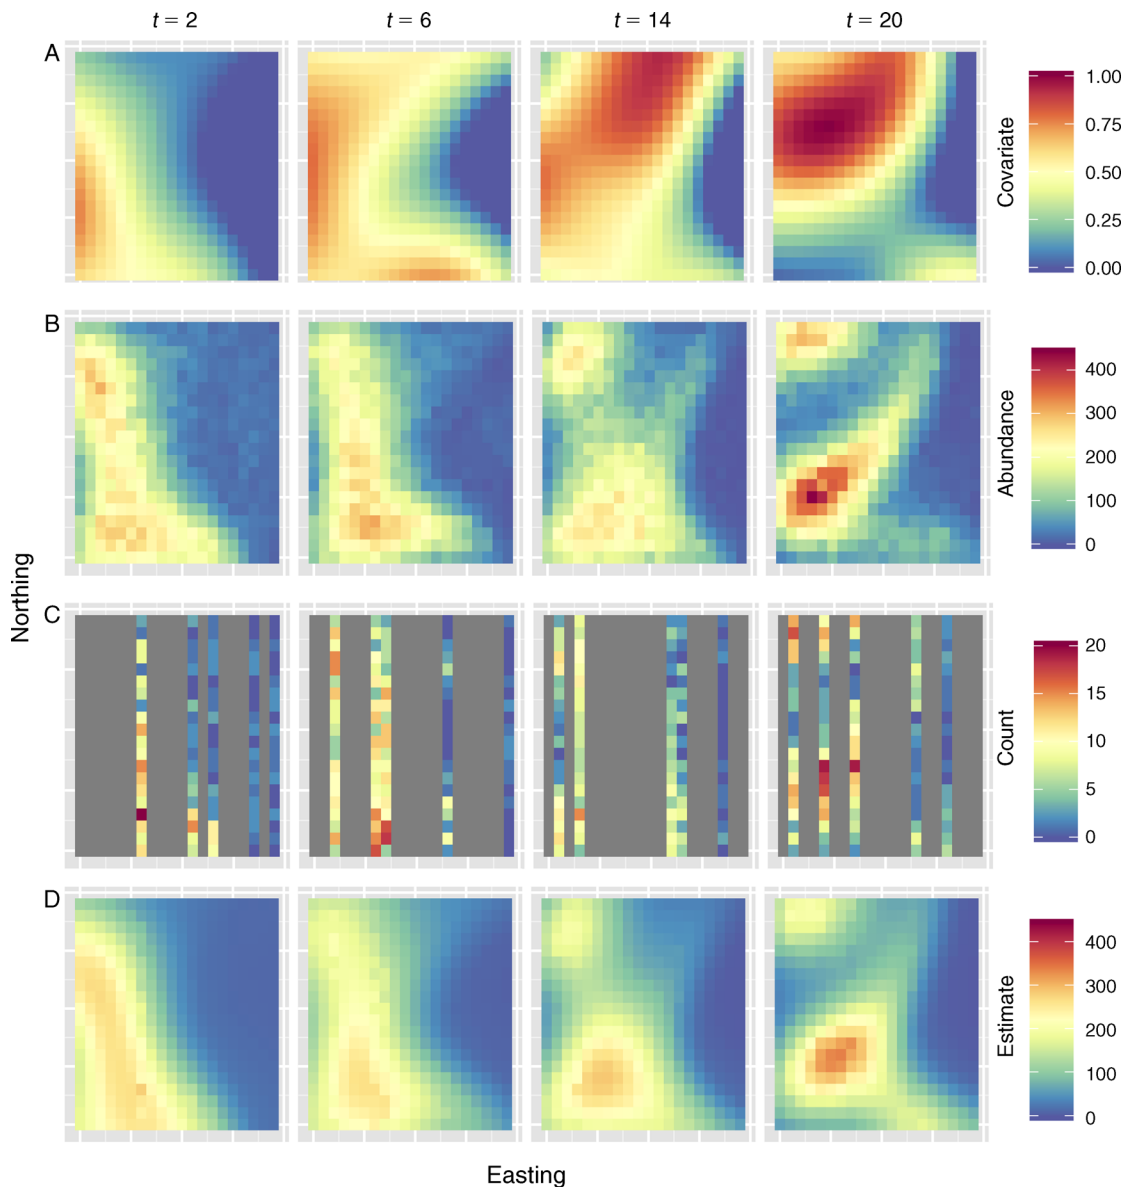


FIG. 2. An example of one simulation iteration, whereby (A) a generic habitat covariate is simulated across a hypothetical landscape and allowed to evolve over time, (B) animal abundance (number of individuals) is simulated conditional on the simulated habitat covariate and an assumed formulation for spatiotemporal dynamics, (C) transect surveys counting individual animals are simulated over the landscape (here, transects were assumed to cover 5% of each grid cell transversed), and (D) animal abundance is estimated via a spatiotemporal statistical model. Each column depicts results for a given time step (t ; each simulation used 20 time steps). In this particular simulation, abundance was simulated according to the closed-population resource selection (CPRS) specification, while estimation was conducted assuming a spatiotemporal process convolution (STPC) model. Each simulation is given as a 20×20 grid; as simulated landscapes were purely hypothetical, the UTM orientation of the plots is completely arbitrary.

$$E \left[\sum_t (\hat{N}_t^{(m)} - \bar{N})^2 \mid \text{Data} \right]$$

where $\hat{N}_t^{(m)}$ is a posterior realization of abundance at time t (e.g., at the m th iteration of the Markov chain) and E is expected value. This approach seeks to minimize Bayes risk with a squared error loss function. For further information on estimator performance calculations, see Appendix B.

Generic simulation results

Results of generic simulations indicated that the posterior mean for animal abundance produced by the four different spatiotemporal estimation models (AST, CPIF, OPRS, and STPC) were relatively unbiased (Fig. 3, Appendix B: Table B2). There was evidence for a small (<3%) positive bias for the OPRS model, while other estimation models had bias at or near zero. Mean squared error favored the STPC model over the other

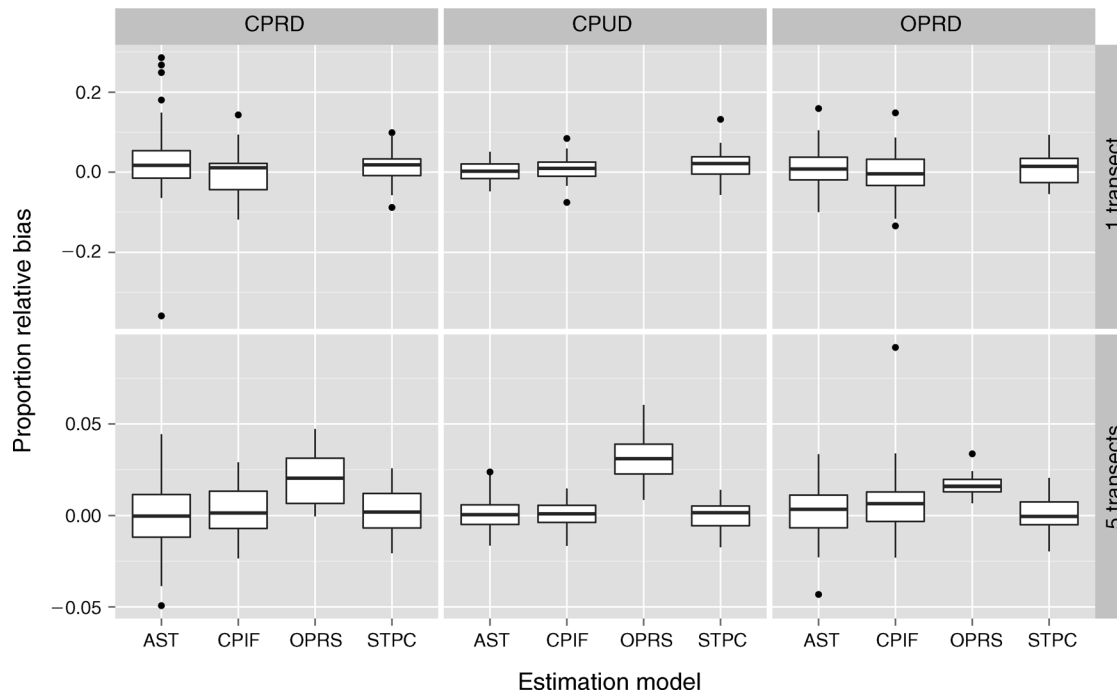


FIG. 3. Boxplots depicting proportion relative bias of animal abundance for different combinations of data-generating model (closed-population model with restricted dispersal [CPRD], closed-population model with unlimited dispersal [CPUD], or an open population model with restricted dispersal [OPRD]), sampling effort (one or five transects/time step), and estimation model (additive space-time [AST], closed-population ideal free [CPIF], open-population resource selection [OPRS], or spatiotemporal process convolution [STPC]) for generic simulations. The lower and upper limits of each box correspond to first and third quartiles, while whiskers extend to the lowest and highest observed bias within 1.5 interquartile range units from the box. Outliers outside of this range are denoted with points. Horizontal lines within boxes denote median bias. Note that the number of simulations for each boxplot varies by estimation model and number of transects (see Appendix B: Table B1).

three estimation models in four out of six design points examined (Appendix B: Table B2). The CV was extremely low for the CPIF estimation model (0.01–0.03; Appendix B: Table B2), which resulted in 90% credible coverage substantially less than nominal (0.18–0.89) for many simulation design points. In fact, coverage was only close to nominal for the CPIF model when a similarly structured model (CPUD) was used to simulate the data. Note that credible interval coverage is nominal whenever the percentage of simulations that include the true abundance value is equal to the stated percentage of the credible interval (i.e., a 90% credible interval should nominally cover true values in 90% of the simulations). The CV for the STPC model was also low (0.02–0.06), leading to coverage levels between 0.34 and 0.70. Precision for the AST model was a little lower (0.03–0.07), resulting in better (but not nominal) coverage levels between 0.73 and 0.86. Coverage rates for the OPRS model were close to nominal (0.85–0.92).

EXAMPLE: SPOTTED SEALS IN THE BERING SEA

We demonstrate the utility of spatiotemporal models for estimating abundance by analyzing count data from spotted seals obtained from aerial transects over the Bering Sea during April and May 2012. Conducted by the U.S. National Marine Fisheries Service (NMFS),

these surveys used thermal imagery to locate seals on ice (see Plate 1); examination of simultaneous high-resolution digital photographs provided information on species composition (for detailed information on survey methods, see Conn et al. 2014). In practice, a variety of observation errors (including incomplete detection) occurred with these data. However, since our main focus is on studying the behavior of different types of spatiotemporal process models, we treat counts as if they were observed perfectly in subsequent modeling (see *Discussion* for further consideration of observation error).

Calibration with simulated data

Before fitting spatiotemporal abundance models to real data, we first conducted a small calibration exercise to examine the performance of three different estimation models (AST, CPIF, STPC) when confronted with distribution and abundance similar to what we expected for spotted seals and the exact levels of sampling effort applied by NMFS in 2012. We did not employ the OPRS estimation model for this exercise, as computational expense was prohibitive. Previous analysis of Bering sea transect counts (Conn et al. 2014, Ver Hoef et al. 2014) indicated spotted seal abundance is maximized in sample units with sea ice proportions in the 0.5–0.8

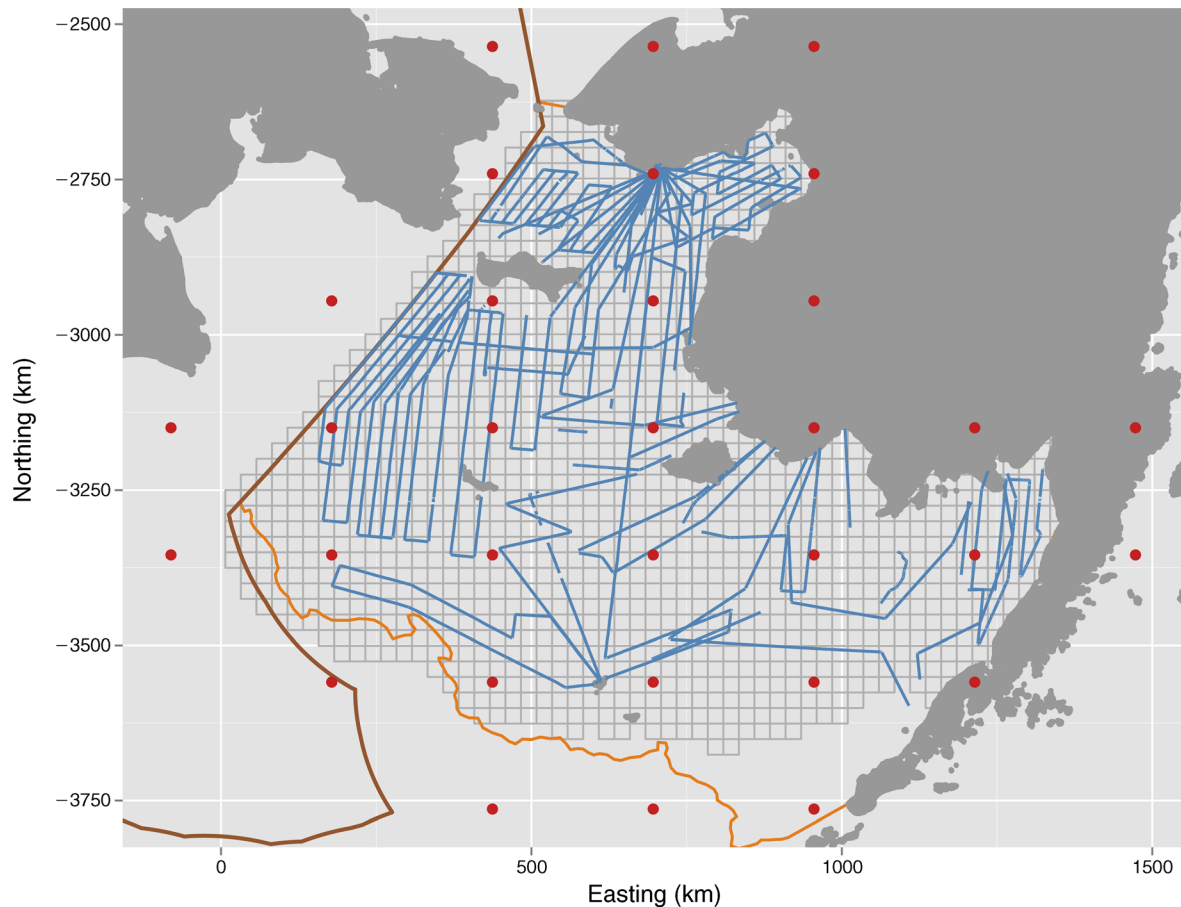


FIG. 4. A map depicting on-effort aerial survey tracks (blue lines) overlaid on the 2012 survey area for ice-associated seals in the Bering Sea. The black-outlined squares are approximately 625 km² and were used to define individual sample units in spatiotemporal analyses. Alaska, USA borders the study area to the west, while Russia appears to the east. Also shown is a 1000-m depth contour (red line), and the United States/Russia Exclusive Economic Zone (EEZ) boundary (orange line), which was used to construct the eastern edge of the survey area. We relied on data from maximal ice extent in April of 2012 to define the southern boundary of the survey area. Red points indicate the position of knots used to construct spatial or spatiotemporal random effects during estimation.

range and near the southern ice edge, and that the abundance of seals in the survey area is roughly constant during the time of surveys. Conditioning on covariate values observed in 2012 (including daily sea ice concentrations, daily distance from southern ice edge, and distance from 1000-m depth contour), we constructed models for spotted seal distribution and abundance (Appendix B). We then used data on the times and locations where transects were flown to sample each virtual population, and each of the three estimation models to conduct posterior inference. Details are provided in Appendix B.

Spotted seal analysis

Next, we analyzed actual counts of spotted seals observed on sea ice during aerial surveys over the eastern Bering Sea between 10 April and 8 May 2012. During this time period, a total of 29 aerial transects were conducted over a 767 114-km² study area (Fig. 4).

An effort was made to distribute sampling effort both spatially and temporally, but field logistics (in particular, weather conditions and snow levels at an airstrip on St. Paul Island, Alaska, USA) precluded flying the desired number of transects along the southern ice edge. The frequency of spotted seal counts was much higher along the southern ice edge (Fig. 4; also see the video in Appendix D), leading to large variation in spotted seal counts as a function of day of survey (Fig. 5). After intersecting our survey effort and observations with a grid of sample units (Fig. 4) and aggregating data by day (i.e., using a daily time step), we attempted to fit three spatiotemporal models to the count data (AST, CPIF, STPC). Details on model configuration and execution are provided in Appendix C.

Spotted seal results

Spotted seal simulations suggested a moderate (3–6%) positive bias for the CPIF and STPC estimation models

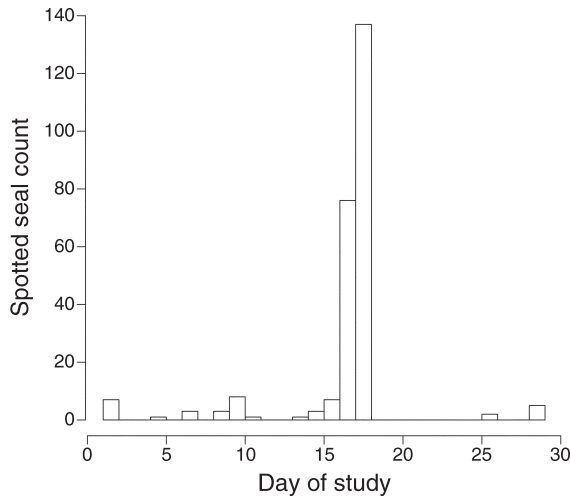


FIG. 5. A histogram depicting spotted seal (*Phoca largha*) counts by day of the study. There was a large spike in counts on days 17 and 18 (26 and 27 April 2012, respectively), owing to flights along the south-central and southeastern ice edge.

(Fig. 6; Appendix B: Table B4), with a large (26%) positive bias for the AST model. Inspection of time series of AST estimates indicated a spurious negative trend in abundance up until days 17 and 18 of the study (the days flights were made into the epicenter of

simulated spotted seal abundance). Restricting estimates of abundance to these dates resulted in much better performance of the AST estimator (Fig. 6; Appendix B: Table B4). Overall, the CPIF estimator had better mean squared error and precision than the AST or STPC estimators, with credible interval coverage close to nominal (Appendix B: Table B4).

Analysis of real-life spotted seal data produced varying results. Despite varying the input configuration of the AST model in several ways, we failed to obtain convergence of this model to a reasonable solution (estimated precision for τ_e converged toward zero so that predicted abundance in unsampled cells was frequently in the millions with many predictions resulting in numerical overflows). The STPC model also produced extremely right-skewed posterior abundance predictions (Appendix C: Fig. C2), and resulted in an apparently spurious increase in abundance (Appendix C: Fig. C3; see *Discussion*).

Compared to the AST and STPC models, the CPIF structure was much more stable, producing spatiotemporal estimates of abundance that mirrored our a priori expectations, with abundance concentrated along the southern ice edge and contracting westward as sea ice diminished in the southeastern quadrant of the study area (Fig. 7; see also the video in Appendix D). Posterior predictions of abundance (Fig. 8) suggested that there

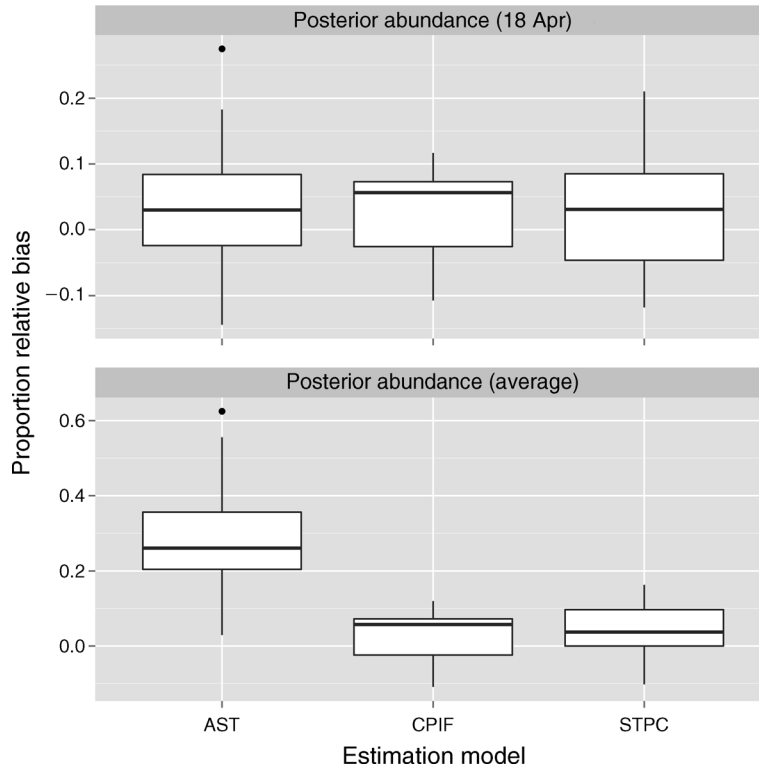


FIG. 6. Boxplots depicting proportion relative bias of animal abundance for different estimation models for spotted seal simulations. Boxplot components are as in Fig. 1. Note that the number of simulations for each boxplot varies by estimation model (100 for AST simulations, 10 for CPIF and STPC). See Fig. 3 for model abbreviations.

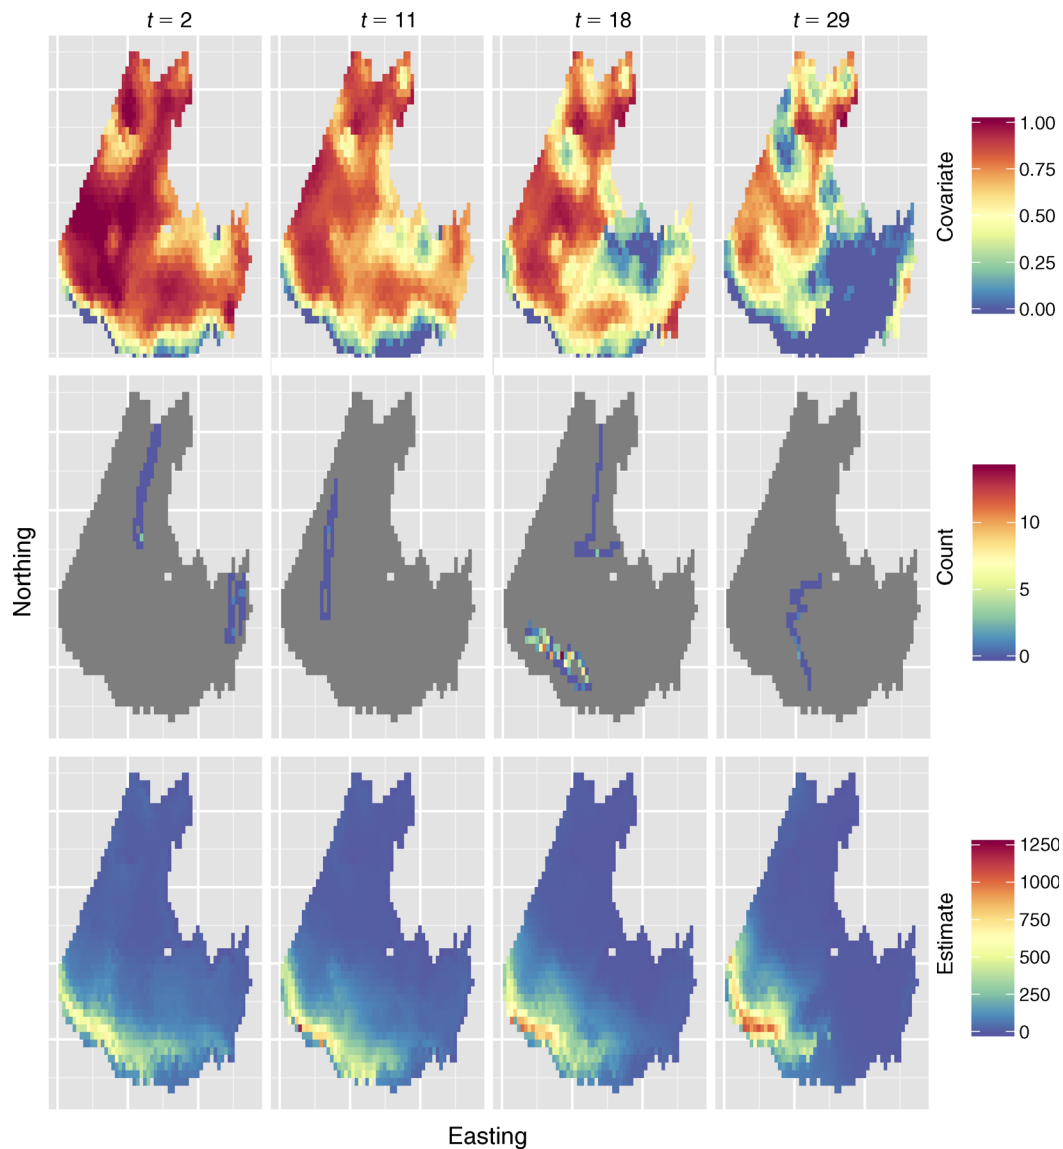


FIG. 7. Time series of sea ice concentration (covariate; top row of panels), survey effort (number of seals counted; middle row), and mean posterior abundance of detectable spotted seals in the Bering Sea (i.e., uncorrected for availability or visibility bias; bottom row). Estimates are made with the CPIF estimation model and are presented only for four selected dates (note $t = 1$ corresponds to 10 April 2012). Each grid cell represents a 25×25 -km area; for information on the spatial extent and geographical location of these estimates, see Fig. 4. For a full time series, see the video in Appendix D.

were approximately 90 000–130 000 detectable spotted seals in the survey area during the 2012 aerial surveys. Examination of posterior distributions for habitat covariates suggested there was a strong quadratic effect of sea ice concentration on spotted seal densities; spotted seal abundance was predicted to be the highest in sample units that were 50% covered by ice, in sample units close to the 1000-m depth contour, and in sample units close to the southern ice edge (Fig. 9).

DISCUSSION

In this study, we examined the performance of several classes of spatiotemporal statistical models when applied

to the specific purpose of estimating animal abundance. Models varied by the underlying structures allowing spatiotemporal variation (i.e., descriptive vs. dynamical), whether populations were assumed open or closed to demographic change, and whether models included spatiotemporal interactions. Several of the models we developed were quite novel; for instance, to our knowledge, this is the first application where a resource selection model has been embedded into a spatiotemporal statistical model for count data (though see Royle et al. [2013] for an example integrating resource selection into capture–recapture models). In the real world, selection of an appropriate statistical model should be

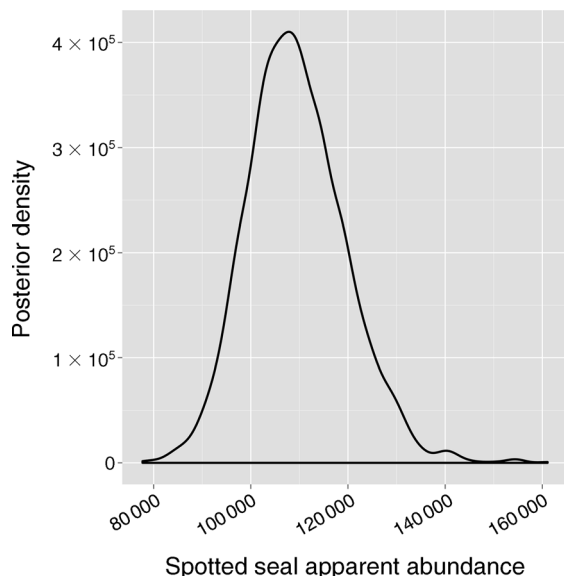


FIG. 8. Posterior predictions of cumulative spotted seal abundance in the eastern Bering Sea between 10 April and 8 May 2012 as produced by the CPIF model. Note that this prediction does not account for incomplete detection.

guided by a number of factors, including the anticipated spatiotemporal dynamics of the study population, the length of the study, and sampling intensity. Selection of an appropriate model should also depend on the performance and robustness of available spatiotemporal models, as inferred from simulation studies and exhibited through case studies.

Results of our generic simulations indicated that spatiotemporal estimation models were capable of producing posterior distributions of animal abundance that were unbiased and highly precise over two different levels of sampling effort and three different data-generating models. These are extremely desirable qualities in an estimation procedure for animal abundance. Unfortunately, credible interval coverage was frequently too narrow (particularly for the CPIF and STPC estimation models), especially when the estimation and data-generating models differed considerably. This is perhaps not surprising, as our estimates are conditional on the assumed models. However, it suggests the need to be careful in taking statements about precision at face value, as model-conditioned standard errors will almost certainly be too small to represent the complex real-world laws governing the distribution and abundance of animal abundance. Two solutions to this predicament may be worth considering. First, one could conduct simulation exercises designed to mimic the real-world system as closely as possible in order to generate an ad hoc variance inflation factor. Second, one could conduct multi-model inference (e.g., using Bayesian model averaging; Hoeting et al. 1999; Hooten and Hobbs 2015) to potentially remove some of the most problematic aspects of conditioning on a single model.

Simulations revealed a large amount of variation in requisite computing times for spatiotemporal statistical models (Appendix B), with the AST model (including descriptive, additive spatiotemporal variation) being much faster than other approaches. The two models allowing spatiotemporal interactions via dynamical process convolution (STPC and CPIF) were somewhat slower, but still allowed us to obtain sufficient MCMC realizations in a reasonable amount of time. The resource-selection model (OPRS), however, took substantially longer, even after expending considerable effort to improve computational speed via a disturbance implementation of the Kalman filter (Durbin and Koopman 2002; see Appendix A). This is likely because the OPRS model requires that one update latent abundance for each sample unit and time step (dimension ST) at every MCMC iteration. By contrast, the AST, STPC, and CPIF specifications often only require updates be made at time and locations where sampling occurs; abundance at unobserved locations can be generated via posterior prediction (see Appendix B). As such, the computational demands required to integrate resource selection into spatiotemporal abundance models may be too onerous for many real-world applications (at present). For instance, trial runs with spotted seal data indicated that execution times would likely be ~ 40 times slower with the OPRS model than STPC or CPIF. Faster run times could have been achieved by modeling resource selection among a smaller number of large sampling units. However, our feeling was that this approach would have degraded the resolution and explanatory power of sea ice covariates.

Simulations specifically tailored to spotted seal counts in the eastern Bering Sea indicated reasonable performance of STPC and CPIF model structures, albeit with a small positive bias. However, the AST formulation often led to a spurious trend in abundance, and estimates that were too large. Estimates from the AST model were much better when confined to 27 April 2012, the day when a transect was flown over the apparent epicenter of spotted seal range. However, fitting these same models to real spotted seal data provided mixed results. The CPIF model produced reasonable model results with estimates that by and large met our a priori expectations (e.g., with spotted seal concentrations along the southern ice edge that gradually contracted and moved northward as ice melted). On the other hand, the AST and STPC formulations, both of which rely on a log-Gaussian Cox process, resulted in extremely right-skewed posterior abundance predictions, to the point where AST results were nonsensical. Our experience (Ver Hoef and Jansen 2007, Conn et al. 2014) is that lack of fit in count data, coupled with sparse sampling, can often cause positive bias in abundance when using log-Gaussian Cox formulations. In effect, lack of fit will tend to result in increased variance on the log of abundance intensity, which translates into positive bias, since abundance effects are modeled in log space. It

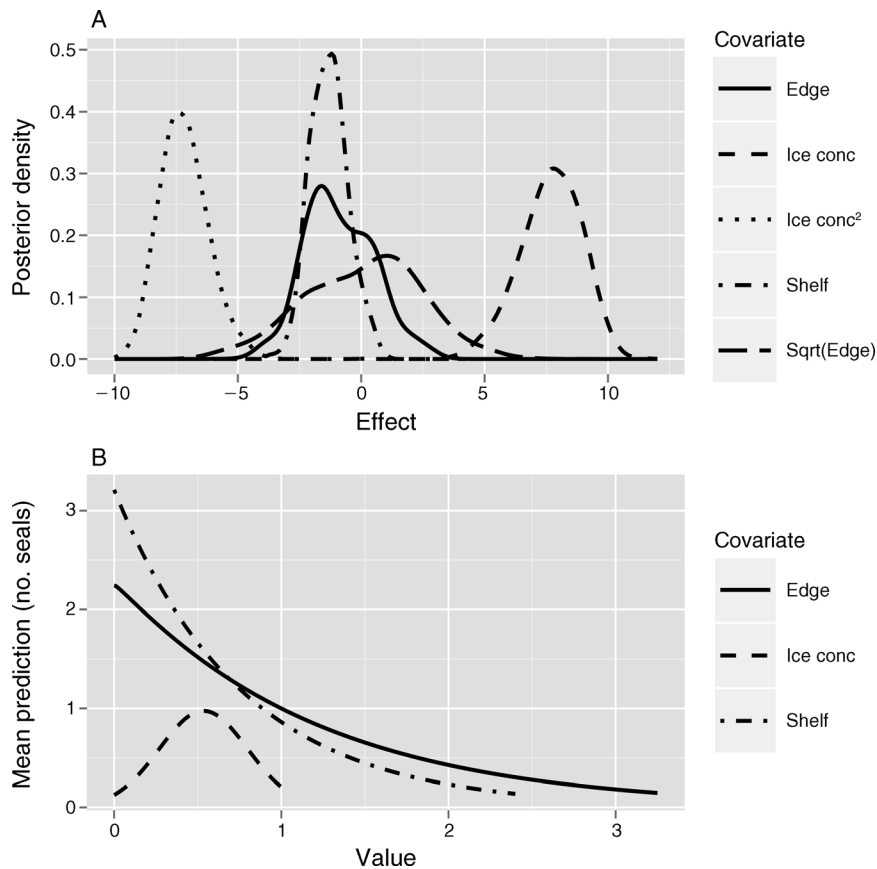


FIG. 9. (A) Posterior distributions and (B) mean relationships for the relationship between covariates and spotted seal abundance (log scale) from CPIF model. Covariates include distance from the southern sea ice edge (Edge), proportion of the sample unit covered with sea ice (Ice conc), and distance from the 1000-m depth contour (Shelf), and included a quadratic effect of ice concentration (Ice conc²) and a square-root effect of distance from sea ice edge (Sqrt(Edge)). Posterior distributions shifted to the left of zero in panel (A) are indicative of negative effects of a covariate on abundance, while shifts to the right are indicative of positive effects on abundance. Panel (B) shows the actual effects of each covariate on relative abundance, calculated using the average value of the other covariates across space and time (e.g., when plotting effects of sea ice concentration, we used the average value of Edge and Shelf). For instance, a sample unit with a sea ice proportion of 0.5 and average values for Edge and Shelf is predicted to have roughly twice the density of spotted seals as would a sea ice proportion of 0.25 or 0.8.

might be possible to decrease lack of fit by employing additional modeling constructs above and beyond the Poisson distribution, such as zero inflation (Agarwal et al. 2002, Ver Hoef and Jansen 2007) or negative binomial models. However, spotted seal data are already sparse, and inclusion of additional structures would require an even greater number of parameters to be estimated. Another possible solution, as suggested by Ver Hoef and Jansen (2007), is to generate a point estimate that minimizes a linex loss function (Varian 1975). However, this puts the analyst in the unenviable position of having a point estimate that is potentially on the far left side of the posterior distribution. By contrast, going with a model such as CPIF presents no such complications.

The models developed and compared in this study are not intended to be definitive regarding the superiority of different types of spatiotemporal models (e.g., descriptive vs. dynamical). Each of the models we implemented

included subjective modeling choices, and it could easily be argued that further work on customizing components of individual models (e.g., by including additional spatiotemporal structure or including such structure at a different stage of the modeling hierarchy) would have influenced relative model performance. Rather, we intend our comparison to help guide ecologists through some of the many modeling decisions and possible frameworks analysts must consider when constructing a spatiotemporal model to estimate animal abundance. That said, there appear to be some real advantages in employing a population closure assumption when it is reasonable to do so. First, it can prevent spurious trends in abundance estimates, particularly near the beginning or end of the time series. Second, one can place a scale prior ($[N] \propto 1/N$) on abundance. This is important, as Link (2013) recently showed that other prior choices sometimes resulted in improper posterior distributions in augmented data applications. We note that the CPIF



PLATE 1. A spotted seal (*Phoca largha*) basking on sea ice in the eastern Bering Sea. The spatial distribution of ice-associated seals can change markedly during spring surveys due to convection and contraction of seasonal sea ice, making them ideal subjects for spatiotemporal statistical modeling efforts. Photo credit: Michael Cameron.

model could be generalized to relax the population closure assumption by modeling abundance at different time steps using a superpopulation formulation (as with the Schwarz-Arnason parameterization of the Jolly-Seber model; cf. Schwarz and Arnason 1996).

For clarity of presentation, we did not attempt to fully model observation errors that are present in the seal transect data. For instance, the count data do not account for all seals because some seals are not available to be detected (i.e., are in the water), thermal imagery misses some seals, and there is some species misclassification. As such, spatiotemporal estimates presented in this study are for demonstration only, as they do not reflect animals unavailable for sampling. Conn et al. (2014) developed observation models to account for these types of detection errors, where the offset term (\mathbf{o} in Eq. 1) was used to handle differences in grid cell size and detection probabilities that are less than one. However, inclusion of species misclassification is more involved and computational demands are considerably greater (Conn et al. 2013). Integrating such observation models with spatiotemporal process models is a subject of current research, and is ultimately what will be needed to produce absolute abundance estimates with realistic measures of precision. We expect that future modeling efforts using an extended data set will also provide more

cohesive estimates, since later surveys (i.e., in 2013) had greater coverage in the southwestern quadrant of our study area.

The methods developed in this study bear some resemblance to the family of so-called N -mixture modeling approaches for plot-based sampling (Royle et al. 2004). For example, Chandler et al. (2011) proposed a class of models for estimating animal density from repeated plot-based samples that allowed the number of animals observed on the study area to change over time; however, their approach assumes that animals are associated with a given sample unit throughout the study, and that temporary emigration out of the sample unit is responsible for variation in counts. As such, their approach is more useful when animals are associated with specific home ranges, as opposed to the case we consider here, where habitat itself can be dynamic and where animals can move across the landscape.

This study represents an initial probing of the effectiveness of different types of spatiotemporal models for abundance estimation. Each of the estimation models considered was configured in what seemed to us to be a sensible manner, but there is certainly future opportunity to investigate these choices in a more rigorous fashion. For instance, the spacing and place-

ment of knots for process convolution could have ramifications for estimator performance. Optimal choices are likely a function of both the quality and quantity of count data available; for a related example in the context of predictive process modeling, see Finley et al. (2009). Similarly, reported computing times could likely be decreased with greater effort at the development stage (for instance, coding some of the more computationally intensive calculations in a lower-level language). Our present investigation focused only on a single type of spatial and spatiotemporal structure (spatial convolution), and others (e.g., GMRFs, predictive process models, geostatistical models) are worth investigating. Other dynamical process formulations are also certainly possible, including cellular automata (Hooten and Wikle 2010) and anisotropic diffusion (Wikle et al. 2001).

ACKNOWLEDGMENTS

A large amount of effort went into collecting seal transect data, and we thank M. Cameron, E. Moreland, and other seal researchers who helped design and execute ice-associated seal surveys. Funding for Bering Sea aerial surveys was provided by the U.S. National Oceanic and Atmospheric Administration and by the U.S. Bureau of Ocean Energy Management (Interagency Agreement M12PG00017). The views and conclusions in this article represent the views of the authors and the U.S. Geological Survey but do not necessarily represent findings or policy of the U.S. National Oceanic and Atmospheric Administration. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government. We thank C. Wikle for helpful discussion; comments by J. Laake, B. McClintock, and anonymous reviewers also helped improve the content and quality of this article. All data collected and research activities described were performed under U.S. National Marine Fisheries Service permit #15126-02.

LITERATURE CITED

- Agarwal, D. K., A. E. Gelfand, and S. Citron-Pousty. 2002. Zero-inflated models with application to spatial count data. *Environmental and Ecological Statistics* 9:341–355.
- Babbitt, K. J., and G. W. Tanner. 2000. Use of temporary wetlands by anurans in a hydrologically modified landscape. *Wetlands* 20:313–322.
- Banerjee, S., B. Carlin, and A. E. Gelfand. 2004. Hierarchical modeling and analysis of spatial data. Chapman & Hall/CRC, Boca Raton, Florida, USA.
- Banerjee, S., A. E. Gelfand, A. O. Finley, and H. Sang. 2008. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society B* 70:825–848.
- Barry, R. P., and J. M. Ver Hoef. 1996. Blackbox kriging: kriging without specifying variogram models. *Journal of Agricultural, Biological, and Environmental Statistics* 1:297–322.
- Berliner, L. 1996. Hierarchical Bayesian time-series models. Pages 15–22 in *Maximum entropy and Bayesian methods*. Kluwer Academic Publishers, Dordrecht, Netherlands.
- Box, G. E. P., G. M. Jenkins, and G. C. Reinsel. 2008. *Time series analysis: forecasting and control*. Fourth edition. Wiley, Hoboken, New Jersey, USA.
- Buckland, S., D. Anderson, K. Burnham, J. Laake, D. Borchers, and L. Thomas. 2001. *Introduction to distance sampling: estimating the abundance of biological populations*. Oxford University Press, Oxford, UK.
- Calder, C. A., C. Holloman, and D. Higdon. 2002. Exploring space-time structure in ozone concentration using a dynamic process convolution model. Pages 165–176 in C. Gatssonis, R. E. Kass, A. Carriquiry, A. Gelman, D. Higdon, D. K. Pauler, and I. Verdinelli, editors. *Case studies in Bayesian statistics 6*. Springer-Verlag, Berlin, Germany.
- Chandler, R. B., J. A. Royle, and D. I. King. 2011. Inference about density and temporary emigration in unmarked populations. *Ecology* 92:1429–1435.
- Cochran, W. 1977. *Sampling techniques*. Third edition. Wiley, New York, New York, USA.
- Conn, P. B. 2014. STabundance. figshare. <http://dx.doi.org/10.6084/m9.figshare.1209622>
- Conn, P. B., B. T. McClintock, M. F. Cameron, D. S. Johnson, E. E. Moreland, and P. L. Boveng. 2013. Accommodating species identification errors in transect surveys. *Ecology* 94:2607–2618.
- Conn, P. B., J. M. Ver Hoef, B. T. McClintock, E. E. Moreland, J. M. London, M. F. Cameron, S. P. Dahle, and P. L. Boveng. 2014. Estimating multispecies abundance using automated detection systems: ice-associated seals in the eastern Bering Sea. *Methods in Ecology and Evolution* 5(12):1280–1293.
- Cressie, N., and C. K. Wikle. 2011. *Statistics for spatio-temporal data*. Wiley, Hoboken, New Jersey, USA.
- Cressie, N. A. C. 1993. *Statistics for spatial data*. Revised edition. Wiley, New York, New York, USA.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B (Methodological)* 39:1–38.
- Durbin, J., and S. J. Koopman. 2002. A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89:603–615.
- Elith, J., and J. R. Leathwick. 2009. Species distribution models: ecological explanation and prediction across space and time. *Annual Review of Ecology, Evolution, and Systematics* 40:677–697.
- Fink, D., W. M. Hochachka, B. Zuckerberg, D. W. Winkler, B. Shaby, M. A. Munson, G. Hooker, M. Riedewald, D. Sheldon, and S. Kelling. 2010. Spatiotemporal exploratory models for broad-scale survey data. *Ecological Applications* 20:2131–2147.
- Finley, A. O., H. Sang, S. Banerjee, and A. E. Gelfand. 2009. Improving the performance of predictive process modeling for large datasets. *Computational Statistics & Data Analysis* 53:2873–2884.
- Fretwell, S. D., and H. L. Lucas. 1970. On territorial behaviour and other factors influencing habitat distribution in birds. *Acta Biotheoretica* 19:16–36.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 2004. *Bayesian data analysis*. Second edition. Chapman & Hall, Boca Raton, Florida, USA.
- Guisan, A., and W. Thuiller. 2005. Predicting species distribution: offering more than simple habitat models. *Ecology Letters* 8:993–1009.
- Haberman, R. 1998. *Elementary applied partial differential equations*. Third edition. Chapman & Hall, Upper Saddle River, New Jersey, USA.
- Hedley, S., and S. Buckland. 2004. Spatial models for line transect sampling. *Journal of Agricultural, Biological, and Environmental Statistics* 9:181–199.
- Higdon, D. 1998. A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics* 5:173–190.
- Higdon, D. 2002. Space and space-time modeling using process convolutions. Pages 37–56 in C. W. Anderson, V. Barnett, P. C. Chatwin, and A. H. El-Shaarawi, editors. *Quantitative methods for current environmental issues*. Springer, London, UK.

- Hoeting, J. A., D. Madigan, A. E. Raftery, and C. T. Volinsky. 1999. Bayesian model averaging: a tutorial. *Statistical Science* 14:382–417.
- Hooten, M., and C. Wikle. 2010. Statistical agent-based models for discrete spatiotemporal systems. *Journal of the American Statistical Association* 105:236–248.
- Hooten, M., C. Wikle, R. Dorazio, and J. Royle. 2007. Hierarchical spatiotemporal matrix models for characterizing invasions. *Biometrics* 63:558–567.
- Hooten, M. B., and N. T. Hobbs. 2015. A guide to Bayesian model selection for ecologists. *Ecological Monographs* 85:3–28.
- Hooten, M. B., and C. K. Wikle. 2008. A hierarchical Bayesian non-linear spatiotemporal model for the spread of invasive species with application to the Eurasian Collared-Dove. *Environmental and Ecological Statistics* 15:59–70.
- Johnson, D. S., J. L. Laake, and J. M. Ver Hoef. 2010. A model-based approach for making ecological inference from distance sampling data. *Biometrics* 66:310–318.
- Johnson, D. S., D. L. Thomas, J. M. Ver Hoef, and A. Christ. 2008. A general framework for the analysis of animal resource selection from telemetry data. *Biometrics* 64:968–976.
- Lele, S. R., and J. L. Keim. 2006. Weighted distributions and estimation of resource selection probability functions. *Ecology* 87:3021–3028.
- Link, W. A. 2013. A cautionary note on the discrete uniform prior for the binomial N . *Ecology* 94:2173–2179.
- Miller, D. L., M. L. Burt, E. A. Rexstad, and L. Thomas. 2013. Spatial models for distance sampling data: recent developments and future directions. *Methods in Ecology and Evolution* 4:1001–1010.
- Miller, M. W., E. V. Pearlstine, R. M. Dorazio, and F. J. Mazzotti. 2011. Occupancy and abundance of wintering birds in a dynamic agricultural landscape. *Journal of Wildlife Management* 75:836–847.
- Moore, J. E., and J. Barlow. 2011. Bayesian state-space model of fin whale abundance trends from a 1991–2008 time series of line-transect surveys in the California Current. *Journal of Applied Ecology* 48:1195–1205.
- Murkin, H. R., E. J. Murkin, and J. P. Ball. 1997. Avian habitat selection and prairie wetland dynamics: a 10-year experiment. *Ecological Applications* 7:1144–1159.
- Patil, G. P., and C. R. Rao. 1978. Weighted distributions and size-biased sampling with applications to wildlife populations and human families. *Biometrics* 34:179–189.
- Prado, R., and M. West. 2010. *Time series: modeling, computation, and inference*. Chapman & Hall/CRC, Boca Raton, Florida, USA.
- R Development Core Team. 2013. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>
- Ross, B. E., M. B. Hooten, and D. N. Koons. 2012. An accessible method for implementing hierarchical models with spatiotemporal abundance data. *PLoS ONE* 7:e49395.
- Royle, J., D. Dawson, and S. Bates. 2004. Modeling abundance effects in distance sampling. *Ecology* 85:1591–1597.
- Royle, J. A., R. B. Chandler, C. C. Sun, and A. K. Fuller. 2013. Integrating resource selection information with spatial capture–recapture. *Methods in Ecology and Evolution* 4:520–530.
- Rue, H., and L. Held. 2005. *Gaussian Markov random fields*. Chapman & Hall/CRC, Boca Raton, Florida, USA.
- Sauer, J. R., and W. A. Link. 2011. Analysis of the North American breeding bird survey using hierarchical models. *Auk* 128:87–98.
- Schwarz, C. J., and A. N. Arnason. 1996. A general methodology for the analysis of capture–recapture experiments in open populations. *Biometrics* 52:860–873.
- Varian, H. R. 1975. A Bayesian approach to real estate assessment. Pages 195–208 *in* S. E. Fienberg and A. Zellner, editors. *Studies in Bayesian econometrics and statistics: in honor of Leonard J. Savage*. North Holland, Amsterdam, Netherlands.
- Ver Hoef, J., and J. Jansen. 2007. Space-time zero-inflated count models of harbor seals. *Environmetrics* 18:697–712.
- Ver Hoef, J. M. 2008. Spatial methods for plot-based sampling of wildlife populations. *Environmental and Ecological Statistics* 15:3–13.
- Ver Hoef, J. M., M. F. Cameron, P. L. Boveng, J. M. London, and E. E. Moreland. 2014. A hierarchical model for abundance of three ice-associated seal species in the eastern Bering Sea. *Statistical Methodology* 17:46–66.
- Wikle, C. 2003. Hierarchical Bayesian methods for predicting the spread of ecological processes. *Ecology* 84:1382–1394.
- Wikle, C. K. 2010. Low rank representations for spatial processes. Pages 89–106 *in* A. Gelfand, P. Diggle, M. Fuentes, and P. Guttorp, editors. *Handbook of spatial statistics*. Chapman & Hall, Upper Saddle River, New Jersey, USA.
- Wikle, C. K., L. M. Berliner, and R. F. Milliff. 2003. Hierarchical Bayesian approach to boundary value problems with stochastic boundary conditions. *Monthly Weather Review* 131:1051–1062.
- Wikle, C. K., and M. B. Hooten. 2010. A general science-based framework for dynamical spatiotemporal models. *TEST* 19:417–451.
- Wikle, C. K., R. F. Milliff, D. Nychka, and L. M. Berliner. 2001. Spatiotemporal hierarchical Bayesian modeling: tropical ocean surface winds. *Journal of the American Statistical Association* 96:382–397.
- Wu, G., S. H. Holan, and C. K. Wikle. 2013. Hierarchical Bayesian spatiotemporal Conway-Maxwell Poisson models with dynamic dispersion. *Journal of Agricultural, Biological, and Environmental Statistics* 18:335–356.

SUPPLEMENTAL MATERIAL

Ecological Archives

Appendices A–D are available online: <http://dx.doi.org/10.1890/14-0959.1.sm>

Data Availability

Data associated with this paper have been deposited in figshare: <http://dx.doi.org/10.6084/m9.figshare.1209622>