

MODELLING OCEAN TEMPERATURES FROM BIO-PROBES UNDER PREFERENTIAL SAMPLING

BY DANIEL DINSDALE AND MATIAS SALIBIAN-BARRERA

University of British Columbia

In the last 25 years there has been an important increase in the amount of data collected from animal-mounted sensors (bio-probes) which are often used to study the animals' behaviour or environment. We focus here on an example of the latter, where the interest is in sea surface temperature (SST), and measurements are taken from sensors mounted on elephant seals in the southern Indian Ocean. We show that standard geostatistical models may not be reliable for this type of data, due to the possibility that the regions visited by the animals may depend on the SST. This phenomenon is known in the literature as preferential sampling, and, if ignored, it may affect the resulting spatial predictions and parameter estimates. Research on this topic has been mostly restricted to stationary sampling locations such as monitoring sites. The main contribution of this manuscript is to extend this methodology to observations obtained by devices that move through the region of interest, as is the case with the tagged seals. More specifically, we propose a flexible framework for inference on preferentially sampled fields where the process that generates the sampling locations is stochastic and moving over time through a two-dimensional space. Our simulation studies confirm that predictions obtained from the preferential sampling model are more reliable when this phenomenon is present, and they compare very well to the standard ones when there is no preferential sampling. Finally, we note that the conclusions of our analysis of the SST data can change considerably when we incorporate preferential sampling in the model.

1. Introduction. The use of animal mounted sensors (bio-probes) to analyse population patterns has grown quickly in the last 25 years [Fedak (2004), Ungar et al. (2005), Evans, Lea and Patterson (2013)], with tags attached to both marine and land based animals. These tags can be used to provide valuable information by collecting data on the environment where the animals live, particularly in regions that are difficult to observe otherwise. One example is given by the use of marine mammal tags to measure oceanographic data such as water temperature, salinity and others. Fedak (2013) highlight the usefulness of such tags in profiling oceanographic data in polar regions where data is typically difficult to obtain.

Although the methodology described in this paper is applicable to a range of different problems, we will focus here on data collected from conductivity-temperature-depth (CTD) bio-probe tags attached to elephant seals in the southern

Received April 2018; revised October 2018.

Key words and phrases. Animal movement models, bio-probes, Laplace approximation, preferential sampling, template model builder.

Indian Ocean. These data were collected and made freely available as part of the Marine Mammals Exploring the Oceans Pole to Pole (MEOP) database [Roquet et al. (2013)], and we utilise the South Indian Ocean data subset described by Roquet et al. (2014). This data set was collected to supplement the Advanced Research and Global Observation Satellite (Argos) float and ship based measures of water masses in typically under sampled areas of the southern oceans, an area which is drastically changing and needs to be further understood [Jacobs (2006)]. While northern oceans have been regularly sampled since the early 2000s using Argo profilers [Gould et al. (2004)], utilising this method in the southern oceans is typically complicated by the presence of sea ice.

The data consist of location coordinates (longitude and latitude which are only available when the animal is surfaced) and corresponding sea surface temperature (SST) measurements. The animal locations in our applied example are determined using Argos and are typically accurate within ± 5 kilometers, while the temperatures are accurate within $\pm 0.03^\circ\text{C}$ [Roquet et al. (2014)]. We use only the temperature data recorded at a depth of 6 meters to represent the SSTs as closely as possible and restrict ourselves to the region between -45 and -65 degrees latitude, between 60 and 120 degrees in longitude over the months of July to September 2012 and use tracks with 50 observations or more. The final data set consists of nine separate tracks with 1630 observations in total, which can be seen in Figure 1.

Geostatistical models and methods [Diggle and Ribeiro (2007)] provide a natural framework to analyse these data. While standard spatial statistical methods consider sampling locations that are chosen independently from the response variable of interest, it is important to note that the movement of the tagged animals




FIG. 1. Plot of all 1630 observations from the subset of southern elephant seal data analysed in Section 5. The first observation was taken on 1 July 2012, and the final observation on 30 September 2012.

(and thus the locations at which measurements are taken) may depend on the variable of interest (SST). Elephant seals are likely to adjust their foraging due to the warming of ocean temperatures [McIntyre et al. (2011)], because their prey tends to prefer cooler and deeper waters. It also appears that there are less successful forages when diving in warmer water [Guinet et al. (2014)].

Preferential sampling [Diggle, Menezes and Su (2010)] refers to the situation where the process that determines the data locations and the spatial field of interest may depend on each other (as can be the case with data collected from animal-mounted tags). The effect of preferential sampling on subsequent inference can be twofold. On the one hand data collected in this way may tend to include a reduced range of the response variable. For example, if seals follow prey towards colder water, their water temperature measurements will tend to not include observations from warmer regions of their environment. On the other hand the locations at which data are collected might carry useful information about the variable of interest. In the previous example one may expect that regions that were not sampled had slightly warmer waters than those visited by the seals. The first issue refers to the fact that the sample of responses may not be a reliable representation over the area of interest, while the second one implies that it may be advantageous for the analysis to take into account the observed locations rather than conditioning on them, as it is done in standard spatial statistical methods.

The impact of preferential sampling on estimation and prediction has been discussed in various recent papers. Shaddick and Zidek (2014) highlighted the preferential nature by which air pollution monitoring sites were added and removed from a UK monitoring network from the 1960s until 2006. Particularly, in the early years of the study, there is evidence that removed sites had a lower annual mean pollution reading than those which were added. Gelfand, Sahu and Holland (2012) showed the prediction effect of preferentially chosen ozone monitoring sites in California, while Diggle, Menezes and Su (2010) and Dinsdale and Salibian-Barrera (2019a) illustrated the effect that preferential sampling of lead concentration in Galicia may have on the resulting predictions. Pati, Reich and Dunson (2011) studied the effect of preferential selection of monitoring sites measuring ozone levels in the eastern U.S.A. More recently, Conn, Thorson and Johnson (2017) showed that preferential sampling in animal population surveys may cause large biases in the animal density estimates, giving an example using aerial survey counts of bearded seals in the eastern Bering Sea.

Research on the issue of preferential sampling has so far mostly been restricted to stationary sampling locations such as monitoring stations. The main contribution of this manuscript is to extend this methodology to observations obtained by devices that move through the region of interest such as those mounted on animals or people. In this paper, the animal locations are modelled using a correlated random walk [Jonsen, Flemming and Myers (2005), Johnson et al. (2008)], where, to allow for the animal movement to depend on the variable of interest, we allow the

drift function (which represents the direction preference at each location) to depend on the SST field. If this field can be assumed to be differentiable, we include a term depending on its gradient to account for the animals' possible preference for particular water temperatures. Finally, we use a behavioural state component as in [Auger-Méthé et al. \(2017\)](#) to allow for a combination of “momentum” and environmental preferences in the animal's movement. Although we believe this movement model to be useful for the analysis of the CTD data, the focus of our work is on studying how to incorporate the possibility of preferential sampling to this type of models rather than advocating for the superiority of a specific movement model over others.

It is interesting to note that the effect of preferential sampling in these “dynamic” spatial models (where sampling locations move through the domain) can be different from what happens when sampling locations are stationary. In the latter case when there is no preferential sampling, observed locations are usually assumed to have a noninformative point pattern sampling distribution over the spatial domain. Thus, when preferential sampling is present it often results in distinct and potentially informative patterns in the monitoring locations [[Diggle, Menezes and Su \(2010\)](#), [Dinsdale and Salibian-Barrera \(2019a\)](#)]. In contrast, even without preferential sampling, the locations visited by the animals in our study would typically not be expected to be evenly distributed over the area of interest (due to the dynamic nature of their movement). As a result the information about SST contained in the locations the animals visited (or in the regions without observations) may be less apparent for these dynamical models than it is for stationary locations. Nevertheless, we show below that not taking into account preferential sampling may still negatively affect the quality and reliability of the estimated parameters and predictions.

To estimate the parameters of interest in the model, we utilise a likelihood approach based on a standard discretisation of the movement equations. The dependence between the sampling locations and the field of interest (SST) results in a likelihood that is computationally challenging to optimise. We follow [Dinsdale and Salibian-Barrera \(2019a\)](#) in using the flexible Template Model Builder R package (TMB) [[Kristensen et al. \(2016\)](#)] to deal with the computational complexity of the preferential sampling analysis.

In many related models discussed in the literature, the underlying field is only assumed to be continuous (but not differentiable). This choice appears to originate on mathematical and computational complexity considerations rather than on the specifics of the phenomenon being studied [[Rue, Martino and Chopin \(2009\)](#), [Lindgren, Rue and Lindström \(2011\)](#), [Simpson, Lindgren and Rue \(2012\)](#)]. Since our model involves the gradient of the underlying SST field [see equation (3.7)], we will consider differentiable Gaussian random fields which are relatively easy to use in the likelihood function when the latter is approximated using using TMB. Details can be found in [Appendix A](#).

Our numerical experiments confirm that better predictions and parameter estimates can be obtained when the model appropriately reflects the potential presence of preferential sampling. Moreover, if the sampled locations are not preferentially chosen, the resulting predictions and parameter estimates are very close to those obtained with the standard model that conditions on the locations. Interestingly, our analysis also shows that the predicted SST in our motivating example obtained with and without a preferential sampling model are in fact different which suggests that the former may be more reliable.

The rest of the paper is organised as follows. Section 2 briefly reviews preferential sampling spatial models and the methods used to perform inference and prediction based on them. Our model for the SST data is discussed in Section 3. The simulations results reported in Section 4 illustrate the benefits of accounting for preferential sampling when it may be present. Section 5 contains the analysis of southern Indian Ocean temperatures from tagged elephant seals, where we compare how accounting for possible preferential sampling might alter the conclusions reached by researchers. A final discussion can be found in Section 6.

2. The preferential sampling problem.

2.1. *Standard model framework.* Our elephant seal data consists of a response variable of interest (SST with possible measurement error, which we will denote by Y), a time stamp and the corresponding location in latitude and longitude coordinates ($\mathbf{X} \in \mathbb{R}^2$), which we assume contains no measurement error; see Section 2.4 for a discussion on this assumption. Since these tags measure water temperature at regular intervals, we consider the temperature measurement taken immediately before a location was obtained (which happens when the animal surfaces). We assume that the measurements taken from each trip are independent from each other.

To model these data we follow the geostatistical framework and notation of Diggle and Ribeiro (2007). More specifically, we assume that the data consist of a finite set of observations from a spatially continuous phenomenon $\{S(\mathbf{x}) : \mathbf{x} \in \mathcal{D} \subseteq \mathbb{R}^2\}$. In our application $S(\mathbf{x})$ denotes the true SST at location \mathbf{x} , and \mathcal{D} is the region of the southern Indian Ocean. Note that $S(\cdot)$ does not vary over time; we discuss this further in Section 2.4. The model for the measurements Y_1, \dots, Y_n obtained in one trip is

$$(2.1) \quad Y_i = \mu + S(\mathbf{X}_i) + Z_i, \quad \mathbf{X}_i \in \mathcal{D}, i = 1, \dots, n,$$

where $\mu \in \mathbb{R}$ is a constant mean parameter over \mathcal{D} and \mathbf{X}_i is the measurement location. The Z_i is included in the model above to account for measurement errors in the sea surface temperatures and are assumed to be mutually independent random variables with mean 0 and so called “nugget variance” τ^2 .

We assume that the SST field S is a Gaussian process with mean 0 and Matérn covariance function given by

$$(2.2) \quad C(r) = \sigma^2 \frac{2^{1-\kappa}}{\Gamma(\kappa)} \left(\frac{r}{\phi}\right)^\kappa K_\kappa\left(\frac{r}{\phi}\right), \quad r > 0,$$

where r is the distance between two points and K_κ is the modified Bessel function of the second kind. The scale (range) ϕ and smoothness κ parameters control the rate of correlation decay over distance and σ^2 the marginal variance of the process (partial sill). Furthermore, κ controls the smoothness of the process realizations. Since our model involves the gradient S (see (3.7)), in what follows we will assume that the surface of SST is differentiable (in mean-square sense) which corresponds to $\kappa > 1$ [Diggle and Ribeiro (2007)]. More specifically, we take $\kappa = 2$ which results in a spatial process S that is mean-square differentiable [Banerjee, Gelfand and Sirmans (2003), Banerjee and Gelfand (2003)].

The main goal of our analysis is to obtain predictions for the SST field S over a grid of unobserved locations. Given estimates for the unknown field parameters in the model one can use standard spatial prediction techniques (e.g., kriging). However, as discussed in Diggle, Menezes and Su (2010) and Dinsdale and Salibian-Barrera (2019a), when the locations \mathbf{X} may depend on the field of interest S , one can obtain better predictions by including the information on S contained in \mathbf{X} . In the rest of this section, we show how to construct an appropriate likelihood function that incorporates the possibility of preferential sampling and how it can be maximised to obtain more accurate parameter estimates and predictions for the field S .

2.2. Maximum likelihood estimation. We use $[\mathbf{A}; \boldsymbol{\theta}]$ to denote the density or probability mass function of the random object \mathbf{A} which depends on a vector of parameters $\boldsymbol{\theta}$. In our case $\boldsymbol{\theta}$ is the vector of all parameters in the model, which can be partitioned as $\boldsymbol{\theta} = (\boldsymbol{\theta}_F^\top, \boldsymbol{\theta}_L^\top)^\top$, where $\boldsymbol{\theta}_F$ are the parameters of the latent field process (e.g., $\boldsymbol{\theta}_F = (\mu, \tau, \kappa, \phi, \sigma^2)^\top$ when using the model from (2.1) and (2.2)), and $\boldsymbol{\theta}_L$ are the parameters of the sampling location distribution. This last vector is divided into $\boldsymbol{\theta}_{L1}$ which are parameters that account for dependence between sampling locations and the latent field, and $\boldsymbol{\theta}_{L2}$ which do not relate to this dependence, $\boldsymbol{\theta}_L = (\boldsymbol{\theta}_{L1}^\top, \boldsymbol{\theta}_{L2}^\top)^\top$.

We then consider the likelihood function $L(\boldsymbol{\theta})$ based on the observed data \mathbf{Y} and \mathbf{X} :

$$(2.3) \quad L(\boldsymbol{\theta}) = [\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta}] = \int [\mathbf{X}, \mathbf{Y}, S; \boldsymbol{\theta}] dS.$$

Typically, one has $[\mathbf{X}, \mathbf{Y}, S; \boldsymbol{\theta}] = [\mathbf{Y}|S, \mathbf{X}; \boldsymbol{\theta}_F][\mathbf{X}|S; \boldsymbol{\theta}_L][S; \boldsymbol{\theta}_F]$. Standard geostastical models assume that the process that selects the measurement locations \mathbf{X} is independent from the response process S (in symbols $[\mathbf{X}|S; \boldsymbol{\theta}_L] = [\mathbf{X}; \boldsymbol{\theta}_{L2}]$), and hence $[\mathbf{X}, \mathbf{Y}, S; \boldsymbol{\theta}] = [\mathbf{Y}|S, \mathbf{X}; \boldsymbol{\theta}_F][\mathbf{X}; \boldsymbol{\theta}_{L2}][S; \boldsymbol{\theta}_F]$. In this case, it follows that $L(\boldsymbol{\theta}) = [\mathbf{X}; \boldsymbol{\theta}_{L2}][\mathbf{Y}|\mathbf{X}; \boldsymbol{\theta}_F]$ and inference about $\boldsymbol{\theta}_F$ can be carried out conditionally on the observed locations \mathbf{X} .

Preferential sampling refers to the situation where the observed locations \mathbf{X} may depend on the unobserved process S . When $[\mathbf{X}|S; \boldsymbol{\theta}_L] \neq [\mathbf{X}; \boldsymbol{\theta}_{L2}]$ care must be taken when constructing the likelihood function in (2.3). In this case, we cannot

simply condition on the sampling locations but should rather use the full likelihood function:

$$(2.4) \quad L(\theta) = \int [\mathbf{X}, \mathbf{Y}, S; \theta] dS = \int [\mathbf{Y}|S, \mathbf{X}; \theta_F][\mathbf{X}|S; \theta_L][S; \theta_F] dS.$$

2.3. *Preferential sampling using template model builder.* Evaluation of the integral (2.4) is computationally challenging and hence optimising the function is a difficult problem. Diggle, Menezes and Su (2010) proposed a Monte Carlo (MC) approximation to a discrete version of the integral, namely

$$(2.5) \quad \int [\mathbf{Y}|S, \mathbf{X}; \theta_F][\mathbf{X}|S; \theta_L][S; \theta_F] dS,$$

where \mathbf{S} is a set of values of S . The exact locations used in the discretisation depends on the model used for the sampling locations and is discussed in more detail in Section 3.2. A direct MC approximation via simulated instances of the vector \mathbf{S} is particularly inefficient, since many of the realisations of \mathbf{S} may not be compatible with the observed measurements \mathbf{Y} . Alternative representations of the likelihood function require sampling from the distribution of the discretised field \mathbf{S} conditional on the observed locations and measurements (i.e., $\mathbf{S}|\mathbf{Y}, \mathbf{X}$) which is generally intractable [Dinsdale and Salibian-Barrera (2019a)]. Other alternatives to inference under preferential sampling have been proposed. For example, Pati, Reich and Dunson (2011) considered a Bayesian alternative and one could also use the R-INLA package [Rue, Martino and Chopin (2009)] which utilises integrated nested Laplace approximation.

A common assumption made in the literature about preferential sampling is that, conditional on the random field S , the sampling locations are static, often modelled via an inhomogeneous Poisson or similar process where the intensity function of $[\mathbf{X}|S; \theta_L]$ depends on S . However, in our application the sampling locations are obtained from a process continuously moving through the two-dimensional domain. For this reason we wish to use a flexible modelling framework in which we can evaluate the likelihood (2.4) efficiently for more complex forms of $[\mathbf{X}|S; \theta_L]$. Although R-INLA provides a highly efficient computational framework, it did not accommodate our relatively complex models for $[\mathbf{X}|S; \theta_L]$ with an underlying smooth process S (e.g., a mean square differentiable SST surface).

We use the R package TMB [Kristensen et al. (2016)] to maximise (2.4) for the dynamic movements models discussed in the next section. This package uses Automatic Differentiation (AD) [Griewank and Walther (2008)] of a Laplace Approximation to the likelihood to efficiently maximise it with respect to the full parameter vector θ . We define the joint negative log-likelihood function

$$(2.6) \quad f(\mathbf{S}, \theta) = -\log([\mathbf{Y}|S, \mathbf{X}; \theta_F][\mathbf{X}|S; \theta_L][S; \theta_F]),$$

and TMB computes an approximation to $\int \exp[-f(\mathbf{S}, \theta)] d\mathbf{S}$, which can be optimised numerically with respect to θ .

The dimension of the integral in (2.5) grows rapidly with the size of the grid that is used to discretise the field S . Important efficiencies can be obtained by using the stochastic partial differential equation (SPDE) approximations for Gaussian fields [Lindgren, Rue and Lindström (2011)] which are also exploited by R-INLA. More specifically, these SPDE approximations allow the use of sparse precision matrices to more efficiently evaluate the high dimensional integral in (2.5). Although the R-INLA package currently only allows continuous but not differentiable fields S in (2.4), it is not difficult to extend the same approach for smoother random fields when using TMB to approximate (2.5). In particular we work with mean-square differentiable random fields [Banerjee, Gelfand and Sirmans (2003), Banerjee and Gelfand (2003)]. Details can be found in Appendix A.

An important goal of this type of analyses is the prediction of SST on nearby locations that were not sampled. To construct predictions that take into account the preferential nature of the data, it is generally not sufficient to use kriging even with parameter estimates obtained through a corrected likelihood function as in (2.4). Such an approach would effectively ignore the dependency between \mathbf{X} and \mathbf{S} [Ferreira and Gamerman (2015), Dinsdale and Salibian-Barrera (2019a)]. Although the true predictive distribution of \mathbf{S} is intractable in most cases when \mathbf{X} and \mathbf{S} are dependent, TMB provides point predictions and prediction variances from the estimated mode of $[\mathbf{S}|\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}]$ at $\boldsymbol{\theta} = \boldsymbol{\theta}_{\text{opt}}$, where $\boldsymbol{\theta}_{\text{opt}}$ is the vector of parameters that maximises (2.5). Specifically, let

$$(2.7) \quad \hat{\mathbf{S}}(\boldsymbol{\theta}) = \underset{\mathbf{S}}{\operatorname{argmin}} f(\mathbf{S}, \boldsymbol{\theta}),$$

then $\hat{\mathbf{S}}(\boldsymbol{\theta}_{\text{opt}})$ is a predictor for the discretised version \mathbf{S} of the process S based on the preferential sampling model (2.4).

2.4. Assumptions. There are two key assumptions made so far in this section that should be noted. First is the assumption that the observed sampling locations \mathbf{X} are the true positions of the animals. In reality there will be some degree of measurement error attached to the sampling locations which depends on the tag type. Argos locations tend to include significant measurement error in comparison to global positioning system (GPS) locations which is common in land-based tracking such as polar bears [Auger-Méthé et al. (2016)] and various birds such as albatrosses [Weimerskirch et al. (2002)] and gannets [Votier et al. (2010)]. However, such systems are less appropriate for marine systems since GPS requires several seconds of exposure to obtain a location estimate [Dujon, Lindstrom and Hays (2014)]. Recently, Fastloc-GPS tags (<http://www.wildtracker.com/Homex.html>) have become more popular due to their improved accuracy compared to Argos and that these systems only require a fraction of a second to obtain a location estimate. When using such data Auger-Méthé et al. (2017) consider the location measurement error negligible enough to be ignored.

For the purpose of this paper, which is to emphasise the preferential sampling problem, we also decided to ignore measurement error in order to provide realistic but not overly complex models. However, this particular compromise between model complexity and computational efficiency by neglecting measurement error may have a serious impact on the analysis in certain situations. In particular large unaccounted sampling location errors may lead to erroneous conclusions on the animal movement and consequently the preferential sampling effect. A possible strategy to incorporate measurement errors in the locations into our model is to include a latent state of true but unobservable locations [see, e.g., [Albertsen et al. \(2015\)](#), [Johnson et al. \(2008\)](#)] which would add further latent states to the integral in (2.3). Note that in this case care will be needed when considering the interplay between S and \mathbf{X} , since both objects will be unobservable.

The second assumption is that although the samples are taken at various time points, that SST depends only on location and not time. Therefore, we can view the continuous SST field S as constant over time. Further research in this area relaxing this assumption would be valuable. Enabling S to vary over time, as it would in real life, would allow for analysis of data over longer periods of time with the model adapting to SST over various seasons and years. In this case, one would consider the data to be of the form $Y_t = S(\mathbf{X}_t, t) + Z_t$ so that $S(\cdot)$ is a function of both location and time. For the real data analysis in Section 5, we consider data across only three months to reduce the impact of a changing temperature field.

3. A preferential movement model. To account for preferential sampling of ocean temperatures, we need to define a model for the location of marine mammals that takes into account possible relationships between movement velocity and ocean temperature. We will use discretised models as functions of the observed locations $\mathbf{X} = (\mathbf{X}(t_1), \dots, \mathbf{X}(t_n))$, in which movement may also depend on previous locations. In these cases, we can write the location at time t_{k+1} as

$$(3.1) \quad \mathbf{X}(t_{k+1}) = g_{\text{NP}}(\mathbf{X}(t_{1:k}), \boldsymbol{\theta}_L) + \varepsilon(\mathbf{t}_{1:k}, \boldsymbol{\theta}_L),$$

where $\mathbf{X}(t_{1:k}) = (\mathbf{X}(t_1), \dots, \mathbf{X}(t_k))$ and $\boldsymbol{\theta}_L$ is now the vector of all movement parameters. The function $g_{\text{NP}}(\cdot)$ is some deterministic movement function where NP stands for nonpreferential and $\varepsilon(\cdot)$ is an error term.

Using the representation in (3.1), under preferential sampling we need to define a movement model of the form

$$(3.2) \quad \mathbf{X}(t_{k+1}) = g_{\text{P}}(\mathbf{X}(t_{1:k}), S, \boldsymbol{\theta}_L) + \varepsilon(\mathbf{t}_{1:k}, \boldsymbol{\theta}_L),$$

where the function g_{P} is now a function of S therefore enabling the movement model to depend on the temperature field.

3.1. A “preferential” CRW model for marine mammal movement. We consider a model similar to the first difference correlated random walk (DCRW) model [Jonsen, Flemming and Myers (2005)]. We wish to include nonregular time intervals to account for possible irregularity in the surfacing of the marine mammals. Maintaining a constant time step through interpolation of the data, as discussed in Hooten et al. (2017b), Jonsen, Flemming and Myers (2005), McClintock et al. (2012) among others, is not possible in our preferential sampling framework. This is due to the necessity of maintaining the link between the sampling locations (\mathbf{X}) and the corresponding latent field measurements (\mathbf{Y}). If we interpolate the trajectory, it is not clear how we would obtain the corresponding Y_i measurements at these interpolated locations other than using a method such as kriging which may dilute any preferential effect that was present in the original data. Another option may be to use thinning [Gurarie et al. (2017)]; however, we wanted to avoid this in our application in this paper due to the limited temporal resolution of the data to which we have access.

We term this model the “preferential correlated random walk” (PCRW) model and assume that the sampling locations $\mathbf{X}(t_1), \dots, \mathbf{X}(t_n)$ follow

$$(3.3) \quad \mathbf{X}(t_{k+1}) = \mathbf{X}(t_k) + \boldsymbol{\mu}(\mathbf{X}(t_{1:k}), S, \boldsymbol{\theta}_L)(t_{k+1} - t_k) + \boldsymbol{\Sigma}(\boldsymbol{\theta}_L) \mathbf{A}_k \sqrt{t_{k+1} - t_k},$$

where \mathbf{A}_k denotes a standard bivariate normal random vector, $\boldsymbol{\Sigma}$ is a 2×2 matrix that corresponds to the variance of the diffusion terms and t_k are the observation times.

To capture various movement patterns such as foraging and directed movement, rather than using discrete behavioural states [Morales et al. (2004), Breed et al. (2009), McClintock et al. (2012)] we propose a continuous behavioural state system similar to Auger-Méthé et al. (2017), Breed et al. (2012). This method was chosen to obtain a differentiable likelihood function through the Laplace approximation outlined in Section 2.3 which would be invalidated with the more commonly used discrete states [Bolker et al. (2013)]. An alternative approach could be to estimate the movement parameters in TMB then follow this with behavioural state estimation using the Viterbi algorithm [Whoriskey et al. (2017)].

The drift function and behavioural states at the measured locations and times satisfy

$$\boldsymbol{\mu}(\mathbf{X}(t_{1:k}), S, \boldsymbol{\theta}_L) = f(\beta_{t_k}) \boldsymbol{\phi}(\mathbf{X}(t_k), S, \boldsymbol{\theta}_L) + (1 - f(\beta_{t_k})) \mathbf{v}(\mathbf{X}(t_{1:k})),$$

where $\beta_{t_k} \in \mathbb{R}$ for all t_k , $\mathbf{v} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ represents the “velocity” of the animal and $\boldsymbol{\phi} : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$ can be thought of as the foraging movement function that depends on the location and the latent temperature field S which models the possible preference of the animals for different water temperatures. The behavioural state function $f : \mathbb{R} \rightarrow [0, 1]$ depends on β_{t_k} and controls the autocorrelation of the movement at each time point. This ensures that when $f(\beta_{t_k}) \approx 1$ then $\boldsymbol{\phi}(\mathbf{X}(t_k), S, \boldsymbol{\theta}_L)$ becomes the expected drift direction, while when $f(\beta_{t_k}) \approx 0$ movement tends in the direction of the current velocity $\mathbf{v}(\mathbf{X}(t_{1:k}))$.

It is important to note that the velocity function may depend on more than just the previous sampling location. Consequently, like the DCRW, our PCRW model is not Markovian. Although one may consider including a latent velocity state (\mathbf{v}) similar to the continuous time correlated random walk model (CTCRW) of Johnson et al. (2008), if this velocity was to depend on the locations \mathbf{X} and also S , such a model becomes drastically more complicated. We discuss this in more detail in Appendix B. However, alternative continuous-time models may also be adaptable to model preferential movement (e.g., the correlated velocity model (CVM) [Gurarie and Ovaskainen (2011)] and functional movement models (FMMs) [Buderman et al. (2016), Hooten and Johnson (2017a)]).

We consider a nonlatent velocity state approximation, taken to be

$$(3.4) \quad \mathbf{v}(\mathbf{X}(t_{1:k})) = \frac{\mathbf{X}(t_k) - \mathbf{X}(t_{k-1})}{t_k - t_{k-1}}$$

and specify a behavioural function

$$(3.5) \quad f(\beta_{t_k}) = \frac{\exp(\beta_{t_k})}{1 + \exp(\beta_{t_k})},$$

so that as β_{t_k} increases, so does the influence of ϕ , whereas when β_{t_k} decreases the current velocity $\mathbf{v}(\mathbf{X}(t_k))$ becomes more of a factor in the movement. Therefore, our PCRW model can be written

$$(3.6) \quad \mu(\mathbf{X}(t_k), S, \theta_L) = \frac{\exp(\beta_{t_k})}{1 + \exp(\beta_{t_k})} \phi(\mathbf{X}(t_k), S, \theta_L) + \frac{1}{1 + \exp(\beta_{t_k})} \mathbf{v}(\mathbf{X}(t_{1:k})),$$

$$\beta_{t_{k+1}} = \beta_{t_k} + \sigma_\beta B_k \sqrt{t_{k+1} - t_k},$$

where B_k are univariate standard normal random variables and $\sigma_\beta > 0$ determines the evolution of the random states β_{t_k} . Note that the inclusion of the random β states means we need to respecify the preferential likelihood for our Laplace approximation, which we show in Appendix C. Initial values for $\mathbf{X}(t_0)$ and β_{t_0} used in our simulation studies are described in Appendix D.1.

Depending on the application of the PCRW model, a variety of forms for the foraging function ϕ may be used. In our case to model the possible tendency of an animal to move towards particular water temperatures when searching for prey we propose the following:

$$(3.7) \quad \phi(\mathbf{X}, S, \theta_L) = -\alpha S(\mathbf{X}) \nabla S(\mathbf{X}),$$

where $\alpha \in \mathbb{R}$, $S(\mathbf{X})$ is the value of the random field at location \mathbf{X} and $\nabla S(\mathbf{X})$ is the gradient of S at \mathbf{X} . Although the parameter α above may appear to be unidentifiable, this is in fact not the case when you consider the full likelihood function which also includes the density functions $[Y|S, \mathbf{X}; \theta_F]$ and $[S; \theta_F]$.

The form of ϕ in (3.7) defines the expected drift as descending (or ascending if $\alpha < 0$) along the gradient of the SST field with a velocity that depends both

on the temperature at the present location and a scalar α . This is somewhat similar to the varying motility surface used by Russell et al. (2018) in an stochastic differential equation (SDE) model to allow the magnitude of the velocity vector to depend on the location. More specifically, we can view the latent field S as a scaled potential surface in which the gradient of S directs the expected movement with a velocity that also depends on the value of S at that location. Potential surfaces have previously been used to model the movement of animals including monk seals [Brillinger, Stewart and Littnan (2008)], elk [Brillinger et al. (2002), Preisler, Ager and Wisdom (2013)] and ants [Russell et al. (2018)] with various estimation methods for potential surface SDE models compared by Gloaguen, Etienne and Le Corff (2018).

It should be noted that appropriate forms of ϕ for particular species may require specialist input, and we do not claim that the one shown in (3.7) is necessarily the best model for seals. The preferential sampling effect may vary over species, locations and possibly even individuals. However, the form in (3.7) may identify preferential movement, and we use it as an example for integrating a foraging function into the preferential sampling framework. Furthermore, in application we may wish to adjust (3.7) to

$$(3.8) \quad \phi(\mathbf{X}, S, \theta_L) = -\alpha S^*(\mathbf{X}) \nabla S(\mathbf{X}),$$

where $S^* = S + c$ for some constant $c \in \mathbb{R}$ specified by the user. We would do this to ensure that S^* is the same sign across the domain, to prevent the switching of movement patterns when S goes from negative to positive or vice versa.

Finally, using (2.1), (3.3)–(3.7) and a finite-differences approximation to the gradient of the field S , we construct a likelihood function that can be maximised numerically, as described in Section 2.3 above. Predicted values for the SST field can be obtained by using (2.7) at the vector of optimal model parameters.

3.2. Relationship between the movement model and likelihood integral. In this section we discuss the form of the discrete grid \mathbf{S} used to approximate the integral in (2.5). The locations at which we need to integrate over depends entirely on the distribution of $\mathbf{X}|S$. Previous literature, in which the sampling locations were point patterns, suited a finely spaced lattice covering the entire domain [Dinsdale and Salibian-Barrera (2019a)]. This was because every point in the domain was a factor in determining the distribution of sampling locations.

In the case of a moving animal measuring a temperature field, this may or may not be required. Take, for example, a model that assumes the animal might be knowingly moving towards distant points of attraction that are related to the latent field, for example, a high prey region with low water temperature. In this case, locations far away from the animal might impact on the movement, hence requiring an approach similar to the point pattern integrals in which we require using a finely spaced lattice covering the entire, or majority, of the domain.

On the other hand the preferential CRW model we are proposing in this paper assumes movement only depends on the animal's immediate vicinity. This can be seen by observing that the only influence of S on the movement is in (3.8) in which the current temperature and gradient of temperature field impacts movement. Hence, it would be more efficient to use a smaller grid for \mathbf{S} which contains only the sampling locations and those areas nearby, which can be used to calculate the temperature gradient at sampled locations using a finite differences approach.

4. Simulation experiments. In this section we discuss the results of a simulation study conducted to illustrate the effect of preferential sampling on the analysis of spatial data where sampling locations are moving through space, such as is the case in our SST data. The goal is to show to what extent incorporating preferential sampling in the model may improve the resulting SST predictions and parameter estimates. Furthermore, our results indicate that when preferential sampling is not present, there is almost no difference between using a model that incorporates preferential sampling and the usual geostatistical model that conditions on the locations. The nonpreferential analysis can be found in Appendix D.3.

We generated 100 data sets, each of them consisting of up to 300 observations on three animal tracks following the Preferential-CRW movement model described in Section 3.1. Data were first generated on a fine time grid, and a subsample was selected to form each of the 100 data sets. Details on the data generation process can be found in Appendix D.1 and the corresponding code found in Dinsdale and Salibian-Barrera (2019b). With our simulation parameters we expect the tracks to oversample cooler regions. Figure 2 shows one simulated data set where the preferential sampling effect is apparent.

4.1. *Parameter estimates.* Note that because of the way the data was generated (by subsampling trajectories created on a relatively fine time scale), the estimated movement parameters may not correspond to their nominal values used to create the data [Gurarie et al. (2017)]. Hence, we report here results for the estimates of the parameters of the spatial process.

Figure 3 shows the boxplots of the 100 estimated parameters for the spatial field process (θ_F) using each of the two likelihoods (standard and accounting for preferential sampling). The grey horizontal lines represent the true values. As expected, when the model does not account for preferential movement, the prevalence of lower temperatures in the sample introduces a negative bias to the estimates for the mean parameter (μ). A similar pattern is observed also for the scale (ϕ) and marginal variance (σ^2). The bias in the variance estimates is likely due to the tracks avoiding higher temperature regions and recording temperatures with a reduced range than they would otherwise. This may also explain the negative bias in the scale estimates. In contrast, using the Preferential-CRW model with TMB results in better parameter estimates. This is particularly noticeable for the estimates of μ . Movement parameter estimates are discussed in Appendix D.2.




FIG. 2. Example of a simulated data set of three tracks generated using the field and movement parameters defined in Section 4 and resulting in a moderate preferential sampling effect.

4.2. *Prediction.* We now turn our attention to the predictions for the underlying spatial field S . The first set of predictions are computed via kriging with parameter estimates obtained from the standard model that conditions on the observed locations, while the preferential sampling ones correspond to the estimated mode of $[\mathbf{S}|\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta}_{\text{opt}}]$, where $\boldsymbol{\theta}_{\text{opt}}$ are the parameter estimates from our preferential likelihood (see equation (2.7)).

Predictions were computed on a 26×26 lattice for each of the $M = 100$ data sets, and we used two different measures of their quality. The root mean square prediction error (RMSPE) over the discrete domain is given by

$$(4.1) \quad \text{RMSPE}_i = \frac{1}{M} \sum_{j=1}^M \sqrt{(\mathbf{S}_{j,i} - \hat{\mathbf{S}}_{j,i})^2},$$

for each location $i = 1, \dots, N = 26^2$ on the prediction grid. Here $\mathbf{S}_{j,i}$ is the true value of field for the j th simulated data set \mathbf{S}_j at the i th prediction location ($\mathbf{S}_j(\mathbf{x}_i)$), and $\hat{\mathbf{S}}_{j,i}$ is the corresponding predicted value.

To compare the resulting predictions while accounting for their variances, we used ignorance scores (IGN) [Roulston and Smith (2002)], which are given by $\text{IGN}(x) = -\log(p(x))$, where p is the predictive density and x is the target forecast which would be the true SST at that location [Gneiting and Raftery (2007), Siegart, Ferro and Stephenson (2014)]. For each of the $j = 1, \dots, M = 100$ preferential simulations we calculated the mean ignorance score (MIGN) of our latent

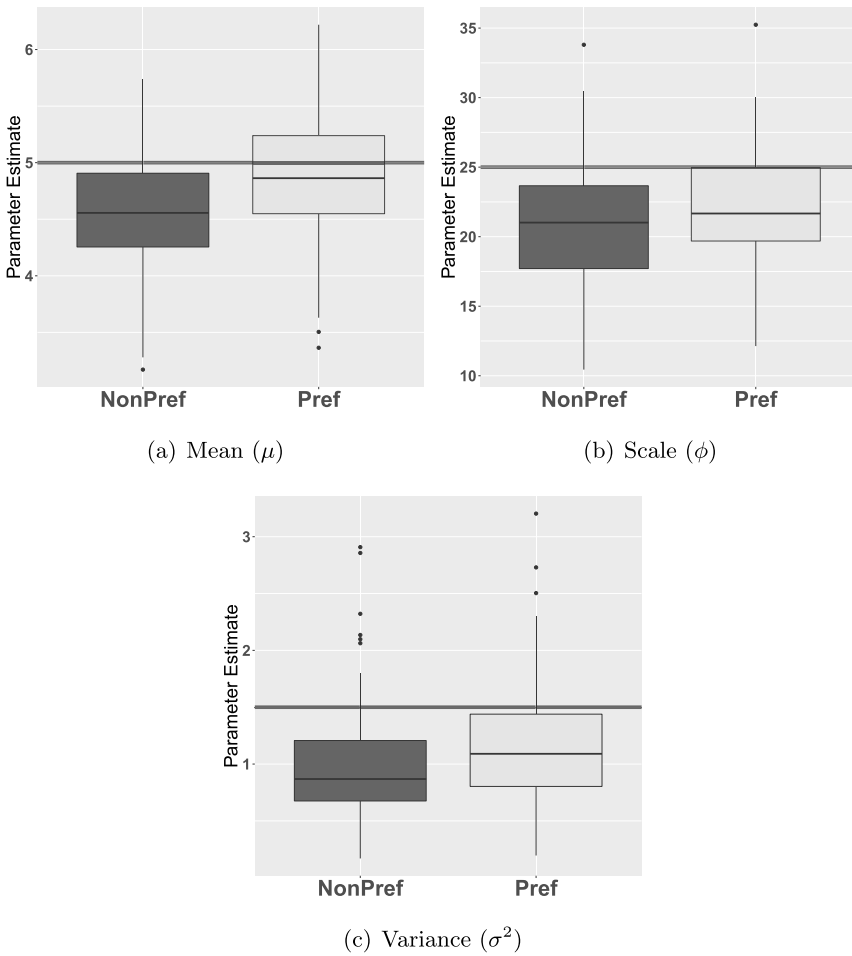


FIG. 3. Field parameter estimates over 100 preferentially sampled simulated data sets with true parameter values marked as grey lines. The abbreviations NonPref and Pref stand for the standard MLE (nonpreferential) estimation and the one using the preferential Preferential-CRW model of Section 3.1.

field predictions as

$$(4.2) \quad \text{MIGN}_j = \frac{1}{N} \sum_{i=1}^N \left\{ \frac{(\mathbf{S}_{j,i} - \hat{\mathbf{S}}_{j,i})^2}{2\hat{\sigma}_{j,i}^2} + \log \hat{\sigma}_{j,i} \right\}, \quad j = 1, \dots, M,$$

where $\hat{\sigma}_{j,i}^2$ is the prediction variance of $\hat{\mathbf{S}}_{j,i}$. This measure gives an indication of model performance for each simulation averaged over the entire domain. We also calculated location-specific ignorance scores, averaging the IGN of each location

over the 100 samples, we call them location ignorance scores (LIGN):

$$(4.3) \quad \text{LIGN}_i = \frac{1}{M} \sum_{j=1}^M \left\{ \frac{(\mathbf{S}_{j,i} - \hat{\mathbf{S}}_{j,i})^2}{2\hat{\sigma}_{j,i}^2} + \log \hat{\sigma}_{j,i} \right\}, \quad i = 1, \dots, N.$$

The LIGN gives an assessment of model prediction across each region of the domain.

To compare the “standard” predictions with the “preferential sampling” ones, we computed the corresponding differences of the above three measures:

$$(4.4) \quad \begin{aligned} \text{RMSPE}_i^{\text{Diff}} &= \text{RMSPE}_i^{\text{P}} - \text{RMSPE}_i^{\text{NP}}, \\ \text{MIGN}_j^{\text{Diff}} &= \text{MIGN}_j^{\text{P}} - \text{MIGN}_j^{\text{NP}}, \\ \text{LIGN}_i^{\text{Diff}} &= \text{LIGN}_i^{\text{P}} - \text{LIGN}_i^{\text{NP}}, \end{aligned}$$

where NP and P indicate the values of the scoring functions for the nonpreferential (standard) and preferential models respectively.

The first panel in Figure 4 shows the values of $\text{RMSPE}_i^{\text{Diff}}$ at each location, colouring the areas in blue for which this measure was negative which correspond to regions where the RMSPEs for preferential sampling predictions were better. Nearly all regions were predicted more accurately on average using the preferential sampling model with only a small number of positive (red) locations.

The boxplot of the $M = 100$ differences in MIGN is displayed in the second panel of Figure 4. Negative values of this difference imply that the P model had a lower average ignorance score (over the region) for that particular simulation run. Again we see that the preferential sampling model performs better than standard methods across the majority of simulated data sets. This conclusion is further supported by the third panel in Figure 4 which shows $\text{LIGN}_i^{\text{Diff}}$ for each lattice prediction location. All locations had a smaller location ignorance score on average when using the preferential sampling model which interestingly includes the locations where point predictions were slightly inferior. This suggests that, even in this case, the prediction variances were more reasonable.

5. Real data example. In this section we analyse southern Indian Ocean temperatures at a depth of 6 meters, which we will be calling SSTs and using CTD sensor data from tags attached to elephant seals. These data were collected and made freely available as part of the MEOP database [Roquet et al. (2013, 2014)] and is described in Section 1 and shown in Figure 1.

Since these data are relatively near to the South Pole, there is a large change in true distance between latitudinal and longitudinal degrees over the domain. For example, we can see in Figure 1 that the change in latitude is not constant in distance between -45° and -65° . This is due to the difficulties of representing locations from a three-dimensional sphere on a two-dimensional surface which is of particular importance near the poles. One option is to use the Haversine formula to calculate the great-circle distance between two points in the latitude/longitude space

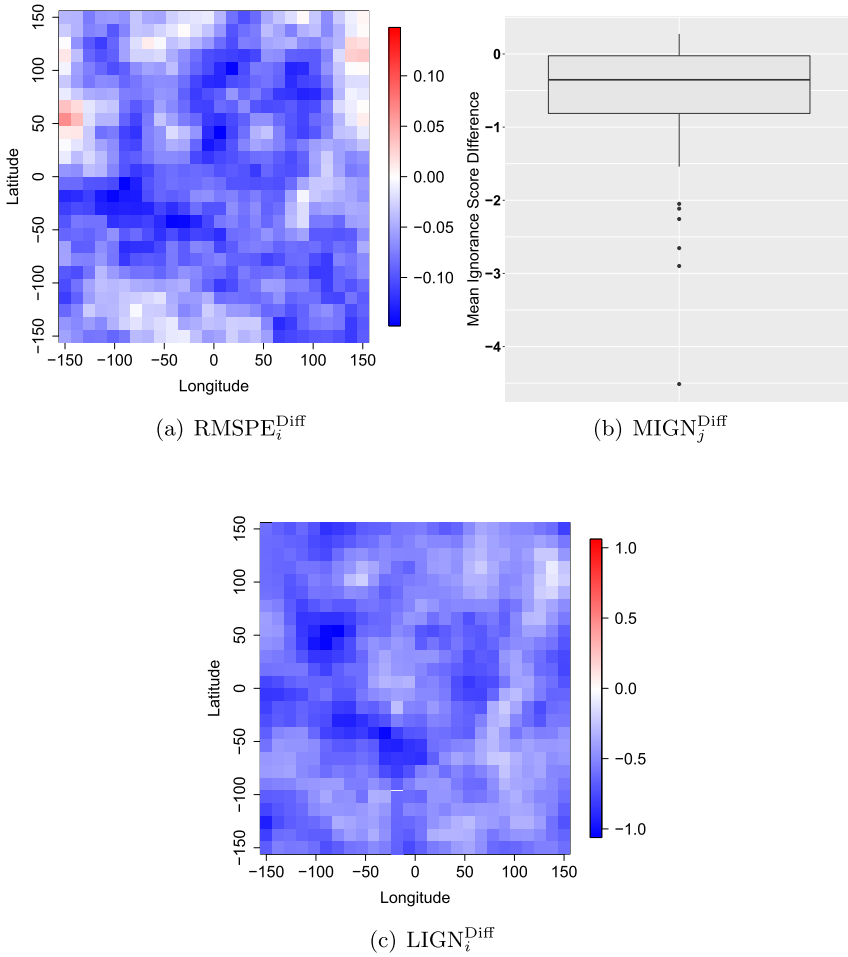


FIG. 4. Comparison of root mean square prediction error (RMSPE) difference, mean ignorance score (MIGN) difference and location ignorance score (LIGN) difference respectively across 100 preferentially sampled simulated data sets.

[Robusto (1957)]. However, as is discussed by Gneiting (2013), Jeong and Jun (2015) among others, using the Matérn covariance along with great circle distances requires the use of a non mean-square differentiable processes ($0 < \kappa \leq 0.5$), since the Matérn class is not isotropic otherwise.

To more accurately represent distance over our domain but remaining in the Euclidean space, we transformed the sampling locations. Our transformation used a scaled version of the Universal Transverse Mercator (UTM) projection (zone 43), which can be seen in Figure 5. This transformation provides us with a more accurate representation of distance than using raw latitudes and longitudes, while remaining in the two-dimensional Euclidean space and retaining the isotropic prop-

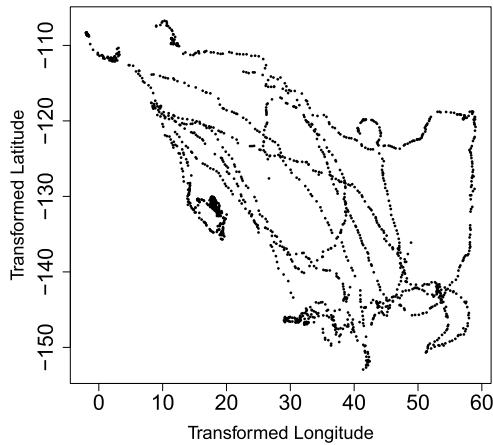


FIG. 5. *The corresponding transformed locations from Figure 1 using a scaled version of the Universal Transverse Mercator (UTM) projection (zone 43).*

erties of our correlation model. The values of our scale have no real-world interpretation, other than to provide more realistic scaled distances between sampled locations.

We compare both the parameter estimates and corresponding field predictions using the PCRW model described in Section 3.1, with those obtained from a standard geostatistical model in which the sampling locations were considered independent of the temperature field. In order to control the computational complexity of the analysis and also to explore the sampling distribution of the field parameter estimators, we use 50 subsamples from the data. For each of these 50 replications we randomly sampled 40 observations from each of the nine tracks and estimated the parameters on the subsampled data. The resulting parameter estimates are displayed in Figure 6. We note that the difference between the two sets of estimates is very small in the scale and variance parameter estimates. However, there was an increase in the mean parameter estimate using the PCRW model. The consistently positive $\hat{\alpha}$ estimates suggest that there was a tendency of the animals to move towards cooler regions according to our model which explains the increased mean parameter estimates.

For predictive assessment of the ocean temperatures obtained from marine mammal tags, we compare our model predictions with monthly average temperature fields obtained via the Simple Ocean Data Assimilation ocean/sea ice reanalysis (SODA) [Carton and Giese (2008)], specifically the SODA version 3 (SODA3) reanalysis [Carton et al. (2018)] which uses all temperature and salinity profiles from the World Ocean Database. This data is available for depths between five and 5000 meters below the surface at a spatial resolution of 0.5 degrees latitude and longitude. Data is available from January 1980 to present.

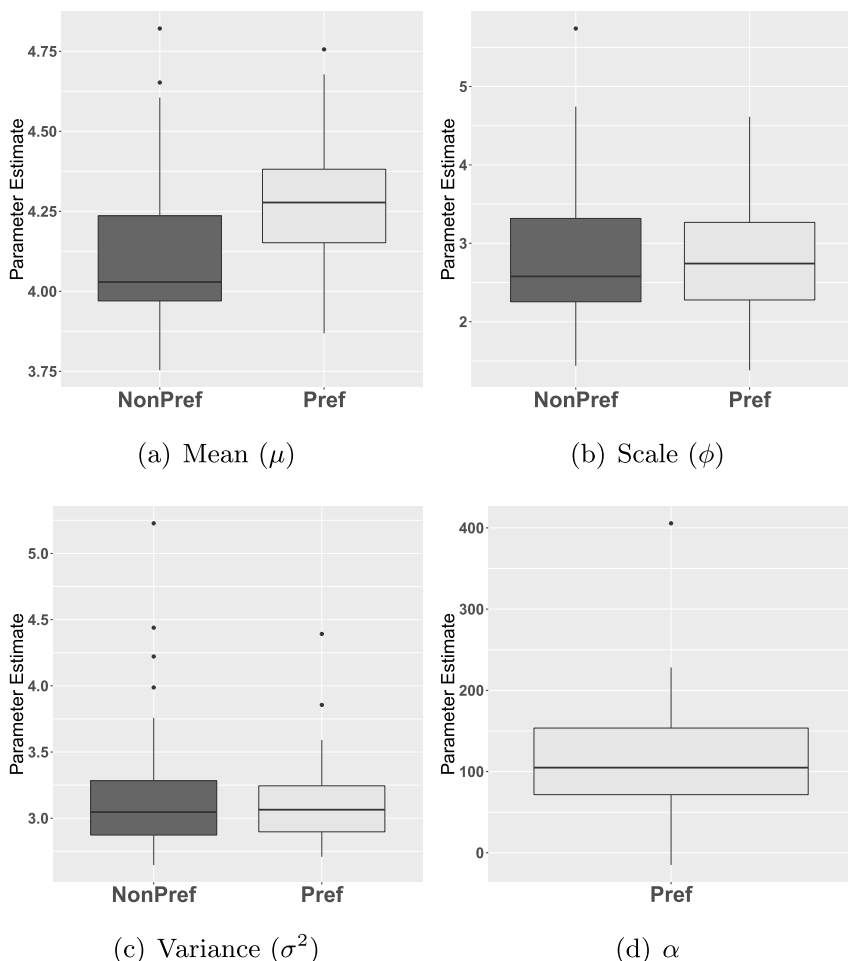


FIG. 6. Comparison between preferential and standard MLEs for the preferential and field parameters over 50 data sets consisting of a subsample of nine tracks with 40 observations each.

The first panel of Figure 7 shows the SODA3 monthly average ocean temperature field for August for a depth of 5 meters. We compare our predictions with these monthly averages at $N = 461$ locations which were chosen as those on the original 26×26 lattice points which were close enough to sampling locations to obtain kriging predictions that did not simply tend to the mean trend parameter $\hat{\mu}$. The locations on the lattice which were not used are shown in grey. To compare point predictions, we consider the quantiles of the difference between the two prediction methods on each of the data sets at each prediction location. In other words we consider the quantiles of each coordinate of the vector $\mathbf{D} = (D_1, \dots, D_N)$ such

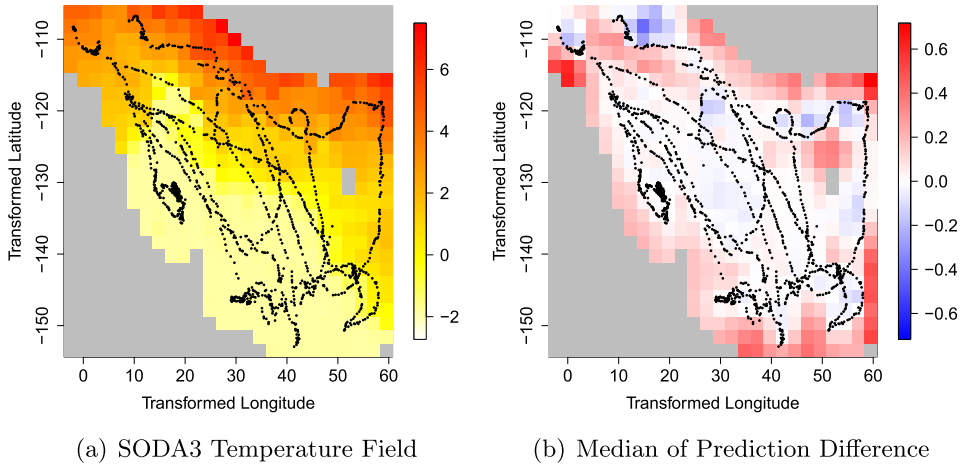


FIG. 7. Panel (a): Monthly average field obtained via the Simple Ocean Data Assimilation ocean/sea ice reanalysis version 3 (SODA3) analysis [Carton *et al.* (2018)]. Panel (b): Median of prediction difference between preferential and nonpreferential models. Positive values imply that the preferential model tends to increase temperature prediction at that location and negative the opposite.

that

$$(5.1) \quad D_i = (\hat{S}_{1,i}^P - \hat{S}_{1,i}^{NP}, \dots, \hat{S}_{50,i}^P - \hat{S}_{50,i}^{NP})^T,$$

for $i = 1, \dots, 461$, where $\hat{S}_{j,i}^P$ is the prediction at location i in simulation j using the preferential model and $\hat{S}_{j,i}^{NP}$ the equivalent using the nonpreferential model. We plot the 50% quantile of \mathbf{D} in the second panel of Figure 7. As we would expect with the positive α estimates, most areas further from sampling locations had an increased SST prediction using the PCRW model. Interestingly however, the PCRW model actually tended to decrease SST prediction in certain areas of the West and North East predictive region. We discuss the magnitude of these prediction differences further in Appendix E.

To assess the prediction accuracy compared to the SODA3 data set, we consider ignorance scores and RMSPEs in the same manner as in Section 4.2. We plot the difference in RMSPE, MIGN and LIGN defined in (4.4) in Figure 8. The first panel shows that the PCRW model tended to reduce RMSPE in comparison to nonpreferential prediction in the northern regions, while underperforming in areas of the South. However, the LIGN in the third panel shows improved ignorance scores in general using the preferential model. An indication of superior prediction performance is through the second panel of Figure 8 which shows the majority of MIGNs across the 50 subsamples were negative and indication of better predictions using the PCRW model.

To summarise, the PCRW model applied using TMB showed differences in mean parameter estimation when compared to the naive method and identified possible

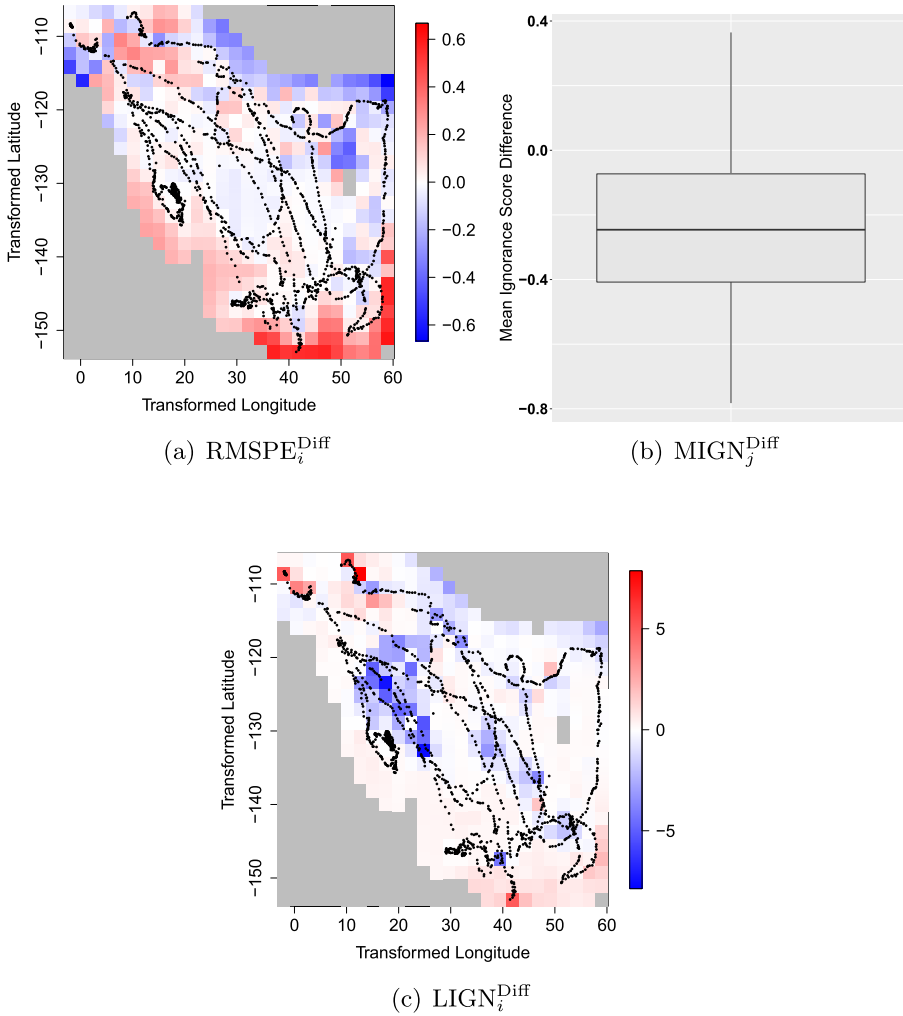


FIG. 8. Comparison of root mean square prediction error (RMSPE) difference, mean ignorance score (MIGN) difference and location ignorance score (LIGN) difference respectively across 50 simulated data sets between preferential and nonpreferential temperature predictions.

tendency of sampler movement towards the cooler regions. This translated to many areas of increased SST prediction when using the preferential model in spatial prediction but not across the entire spatial region. Small changes in SST prediction as shown by our model may be of increased importance when dealing with complex systems as ocean temperatures. Although this simulation used a simplistic relationship between the sampler and the process being modelled, our results suggest that developing more realistic models for this type of data can improve the resulting statistical inference about our ocean climates.

6. Discussion. We have shown the damaging effect that preferential sampling may have on statistical inferences based on spatial models where monitoring locations are not stationary. The simulation experiments reported in Section 4 illustrate how predictions may be improved when accounting for the preferential nature of movement in the sampling model. This is evident in the parameter estimates and also in the predictive performance. Here the combination of corrected parameters and a predictive distribution that accounts for the relationship between the sampling locations and the spatial field of interest improves notably upon the simple extrapolation of kriging. When we compare the results of the standard method (which conditions on the observed locations) and the Preferential-CRW model in our motivating example, we observe increased mean parameter estimates for the SST field while the predicted SST fields show consistent differences.

These analyses highlight the importance of expanding preferential sampling methods to the case of moving sampling locations which is becoming a more prevalent situation. We show how implementing a Laplace approximation to the likelihood function via the R package `TMB` allows for flexible movement model specification. This method can expand beyond the typical point process models to sampling processes derived from movement models which may depend on our field of interest S . This is just the first step in incorporating preferential sampling into the statistical analyses of tagged marine mammals. Considering the observed measurement locations to contain some measurement error by assuming there is an latent true location state would help to account for the uncertainty in the Argos location estimates, while there are many other applications such as land animals and other moving sampling processes as the next step in expanding our applications and methodology. Furthermore, a natural extension of this research is to consider data at any depth. Although SST analysis is a useful first step, an important application of these methods would be to aid in mapping the water masses at depths unobservable by satellites. Theoretically, little would change in comparison to the SST data we consider here in analyses at specific depths; however, the location accuracy underwater is reduced. In this case, location estimates could be inferred based on accelerometer data using methods such as the dead reckoning algorithm [Wilson and Wilson (1988), Liu et al. (2016)] and ocean temperature at specific depths analysed to build a water temperature profile.

APPENDIX A: USING SMOOTH RANDOM FIELDS WITH TMB

Many geostatistical models proposed to analyse spatial data assume that the underlying random field is continuous (but not necessarily differentiable). In what follows we show how `TMB` can be used to efficiently employ mean-square differentiable Gaussian fields (GFs) in our models, using the approach by Lindgren, Rue and Lindström (2011) who show how Gaussian Markov random field (GMRF) representations of GFs with Matérn covariance (as defined in (2.2)), can be constructed through the solution to a stochastic partial differential equation (SPDE) when driven by white noise.

Methods that utilise GMRF representations of GFs require constructing a sparse precision matrix \mathbf{Q} that closely represents the covariance of the GF. We follow directly after (9) in Lindgren, Rue and Lindström (2011) but note the change in notation from κ in the original work to ϕ here and that in our two-dimensional application $\alpha = \kappa + 1$ where κ is the smoothness parameter in (2.2).

Lindgren, Rue and Lindström (2011) write the precision matrix $\mathbf{Q}_{\alpha,\phi}$ as a combination of $m \times m$ matrices \mathbf{C} , \mathbf{G} and \mathbf{K}_ϕ ,

$$\begin{aligned} \mathbf{C}_{i,j} &= \langle \psi_i, \psi_j \rangle, \\ \mathbf{G}_{i,j} &= \langle \Delta \psi_i, \Delta \psi_j \rangle, \\ (\mathbf{K}_\phi)_{i,j} &= \phi^{-2} \mathbf{C}_{i,j} + \mathbf{G}_{i,j}, \end{aligned}$$

where m is the number of vertices in the triangulation of the domain.

This combination depends on α but can be calculated recursively as

$$\begin{aligned} \mathbf{Q}_{1,\phi} &= \mathbf{K}_\phi, \\ \mathbf{Q}_{2,\phi} &= \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{K}_\phi, \\ \mathbf{Q}_{\alpha,\phi} &= \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{Q}_{\alpha-2,\phi} \mathbf{C}^{-1} \mathbf{K}_\phi, \quad \text{for } \alpha = 3, 4, \dots \end{aligned}$$

Unfortunately, \mathbf{C}^{-1} is dense, but Lindgren, Rue and Lindström (2011) show that \mathbf{C} can be replaced by the diagonal matrix $\tilde{\mathbf{C}}$, where $\tilde{C}_{i,i} = \langle \psi_i, 1 \rangle$. Hence, $\tilde{\mathbf{C}}$ is sparse and the resulting precision matrix $\mathbf{Q}_{\alpha,\phi}$ also sparse. Further details can be found in Lindgren and Rue (2015), Lindgren, Rue and Lindström (2011), Simpson, Lindgren and Rue (2012).

A.1. Implementation of $\kappa = 2$ in TMB. R-INLA only implements the SPDE approximation for $\alpha \in (0, 2]$. At the largest smoothness of $\alpha = 2$ ($\kappa = 1$), the corresponding Matérn field is mean square continuous but not mean square differentiable. In some applications this might not suit the problem at hand (e.g., when derivatives of the spatial field may be part of the model). Below we extend the current approximations to the case $\alpha = 3$ ($\kappa = 2$) and show how this can be implemented in TMB via the built-in compatibility with R-INLA.

To begin, note that for $\alpha = 2$ we can expand the formula for $\mathbf{Q}_{2,\phi}$ above as follows:

$$\begin{aligned} \mathbf{Q}_{2,\phi} &= \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{K}_\phi \\ &= (\phi^{-2} \mathbf{C} + \mathbf{G}) \mathbf{C}^{-1} (\phi^{-2} \mathbf{C} + \mathbf{G}) \\ &= \phi^{-4} \mathbf{C} + 2\phi^{-2} \mathbf{G} + \mathbf{G} \mathbf{C}^{-1} \mathbf{G} \\ &= \phi^{-4} \mathbf{M}_0 + 2\phi^{-2} \mathbf{M}_1 + \mathbf{M}_2, \end{aligned}$$

where $\mathbf{M}_0 = \mathbf{C}$, $\mathbf{M}_1 = \mathbf{G}$ and $\mathbf{M}_2 = \mathbf{G}\mathbf{C}^{-1}\mathbf{G}$. Note that the matrices \mathbf{M}_1 , \mathbf{M}_2 and \mathbf{M}_3 do not depend on ϕ or κ and can be computed with the R-INLA function `inla.spde2.matern`. Now we expand $\mathbf{Q}_{3,\phi}$:

$$\begin{aligned} \mathbf{Q}_{3,\phi} &= \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{Q}_{1,\phi} \mathbf{C}^{-1} \mathbf{K}_\phi \\ &= \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{K}_\phi \mathbf{C}^{-1} \mathbf{K}_\phi \\ &= (\phi^{-2}\mathbf{C} + \mathbf{G})\mathbf{C}^{-1}(\phi^{-2}\mathbf{C} + \mathbf{G})\mathbf{C}^{-1}(\phi^{-2}\mathbf{C} + \mathbf{G}) \\ &= \phi^{-6}\mathbf{C} + 3\phi^{-4}\mathbf{G} + 3\phi^{-2}\mathbf{G}\mathbf{C}^{-1}\mathbf{G} + \mathbf{G}\mathbf{C}^{-1}\mathbf{G}\mathbf{C}^{-1}\mathbf{G} \\ &= \phi^{-6}\mathbf{M}_0 + 3\phi^{-4}\mathbf{M}_1 + 3\phi^{-2}\mathbf{M}_2 + \mathbf{M}_2\mathbf{M}_0^{-1}\mathbf{M}_1. \end{aligned}$$

Hence, we can use \mathbf{M}_0 , \mathbf{M}_1 and \mathbf{M}_2 provided by `inla.spde2.matern` to construct $\mathbf{Q}_{3,\phi}$ for a solution to the SPDE approximation for $\alpha = 3$ ($\kappa = 2$). We simply need to take a new combination of these matrices which in TMB can be done easily within the function that computes the likelihood.

APPENDIX B: PREFERENTIAL-CRW MODELS WITH A FIELD-DEPENDENT VELOCITY TERM

If one considers a velocity term \mathbf{v} defined in (3.4) that may depend on both the locations \mathbf{X} and the underlying field of interest S , the likelihood quickly becomes intractable. Considering a latent velocity state in the model, the full likelihood is

$$\begin{aligned} (\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta}) &= \int \int [\mathbf{X}, \mathbf{Y}, \mathbf{S}, \mathbf{v}; \boldsymbol{\theta}] d\mathbf{S} d\mathbf{v} \\ &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \mathbf{v}; \boldsymbol{\theta}_F][\mathbf{X}, \mathbf{S}, \mathbf{v}; \boldsymbol{\theta}] d\mathbf{S} d\mathbf{v} \\ (B.1) \quad &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \mathbf{v}; \boldsymbol{\theta}_F][\mathbf{X}|\mathbf{S}, \mathbf{v}; \boldsymbol{\theta}_L][\mathbf{S}, \mathbf{v}; \boldsymbol{\theta}] d\mathbf{S} d\mathbf{v} \\ &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \mathbf{v}; \boldsymbol{\theta}_F][\mathbf{X}|\mathbf{S}, \mathbf{v}; \boldsymbol{\theta}_L][\mathbf{v}|\mathbf{S}; \boldsymbol{\theta}_L][\mathbf{S}; \boldsymbol{\theta}_F] d\mathbf{S} d\mathbf{v}. \end{aligned}$$

Notice that this factorisation includes the term $[\mathbf{v}|\mathbf{S}; \boldsymbol{\theta}_L]$. If the velocity was to only depend on the latent field \mathbf{S} and not on the location \mathbf{X} , then we could evaluate the likelihood. However, this makes little biological sense in practice. Preferential movement induced by a latent velocity would require this velocity to depend on the field \mathbf{S} at the current location \mathbf{X} of the animal. In other words the velocity would be dependent not just on the latent field but the current location within this field. However, this means that we would only have a tractable form for $[\mathbf{v}|\mathbf{S}, \mathbf{X}; \boldsymbol{\theta}_L]$ and would therefore need to evaluate

$$\begin{aligned} [\mathbf{v}|\mathbf{S}; \boldsymbol{\theta}_L] &= \int [\mathbf{X}, \mathbf{v}|\mathbf{S}; \boldsymbol{\theta}_L] d\mathbf{X} \\ &= \int [\mathbf{v}|\mathbf{S}, \mathbf{X}; \boldsymbol{\theta}_L][\mathbf{X}|\mathbf{S}; \boldsymbol{\theta}_L] d\mathbf{X} \end{aligned}$$

which is difficult to compute due to the complexity of $[\mathbf{X}|\mathbf{S}; \boldsymbol{\theta}_L]$. The factorisation of the form in (B.1) appears intractable and therefore the Preferential-CRW model in which velocities depend on location and the latent field become difficult to implement using likelihood methods when preferential sampling may be present.

APPENDIX C: PREFERENTIAL CRW LIKELIHOOD

In Section 2.3 we outlined how to evaluate the preferential model using a Laplace approximation to the likelihood function. In previous cases this factorisation was of the form $[\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta}] = \int [\mathbf{X}, \mathbf{Y}, \mathbf{S}; \boldsymbol{\theta}] d\mathbf{S}$, however in the case of the Preferential-CRW model we have a second latent vector in the behavioural states $\boldsymbol{\beta} = (\beta_{t_1}, \dots, \beta_{t_n})$. Therefore we need to re-specify the full likelihood for our Laplace approximation routine:

$$\begin{aligned}
 [\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta}] &= \int \int [\mathbf{X}, \mathbf{Y}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}] d\mathbf{S} d\boldsymbol{\beta} \\
 &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_F][\mathbf{X}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}] d\mathbf{S} d\boldsymbol{\beta} \\
 \text{(C.1)} \quad &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_F][\mathbf{X}|\mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_L][\mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}] d\mathbf{S} d\boldsymbol{\beta} \\
 &= \int \int [\mathbf{Y}|\mathbf{X}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_F][\mathbf{X}|\mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_L][\boldsymbol{\beta}|\mathbf{S}; \boldsymbol{\theta}_L][\mathbf{S}; \boldsymbol{\theta}_F] d\mathbf{S} d\boldsymbol{\beta}.
 \end{aligned}$$

Notice that in the Preferential-CRW model we have $[\boldsymbol{\beta}|\mathbf{S}; \boldsymbol{\theta}_L] = [\boldsymbol{\beta}; \boldsymbol{\theta}_{L2}]$ so therefore in the Laplace approximation implemented using TMB we will need to redefine the joint negative log-likelihood as

$$-\log([\mathbf{X}, \mathbf{Y}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}]) = -\log([\mathbf{Y}|\mathbf{X}, \mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_F][\mathbf{X}|\mathbf{S}, \boldsymbol{\beta}; \boldsymbol{\theta}_L][\boldsymbol{\beta}; \boldsymbol{\theta}_{L2}][\mathbf{S}; \boldsymbol{\theta}_F]).$$

APPENDIX D: SIMULATION DETAILS

D.1. Data generation. To generate each track, we initialised a starting location chosen uniformly at random over the two-dimensional domain $[-150, 150] \times [-150, 150]$, then simulated 360 observations with the time between consecutive observations following an exponential distribution with rate parameter $\lambda = 10$. The first 60 positions of each animal were considered a burn-in period and discarded, resulting in 300 remaining observations. Finally, the track was thinned by taking every third observation, to retain a final track of 100 observations.

For each data set we simulated a random field S over the domain and discretised it on a 51×51 grid. We used a Matérn covariance function as in (2.2) with smoothness parameter $\kappa = 2$, scale $\phi = 25$, marginal variance $\sigma^2 = 1.5$ and a constant mean $\mu = 5$. Since we know the true field in generating the tracks, the gradient of the field used to direct movement was approximated using finite differences from points not necessarily on the grid. For the movement model we set $\sigma_\beta = 0.1$, $\boldsymbol{\Sigma} = 3\mathbf{I}_2$, where \mathbf{I}_2 denotes the 2×2 identity matrix, $\alpha = 100$ and ini-

tiated the behavioural states at $\beta_0 = -1.5$. From (3.6) we see that this choice initialises primarily directed movement ($f(\beta_0) = 0.18$), but with a slight influence of the foraging (preferential) function $\phi(\cdot)$. We assume that τ^2 is a known parameter due to the assumption that the measuring device will have a known sampling error and as it is commonly done in the literature, we also assume that the smoothness parameter κ is known.

D.2. Movement parameter estimates. For the preferentially generated data, Figure 9 shows that the movement parameters estimates, with the positive $\hat{\alpha}$ estimates accounting for the tendency of the sampler to avoid warmer warmers, which explains the correction to the mean parameter estimates observed in Figure 3. Although these estimates cannot in general be compared with the values used to

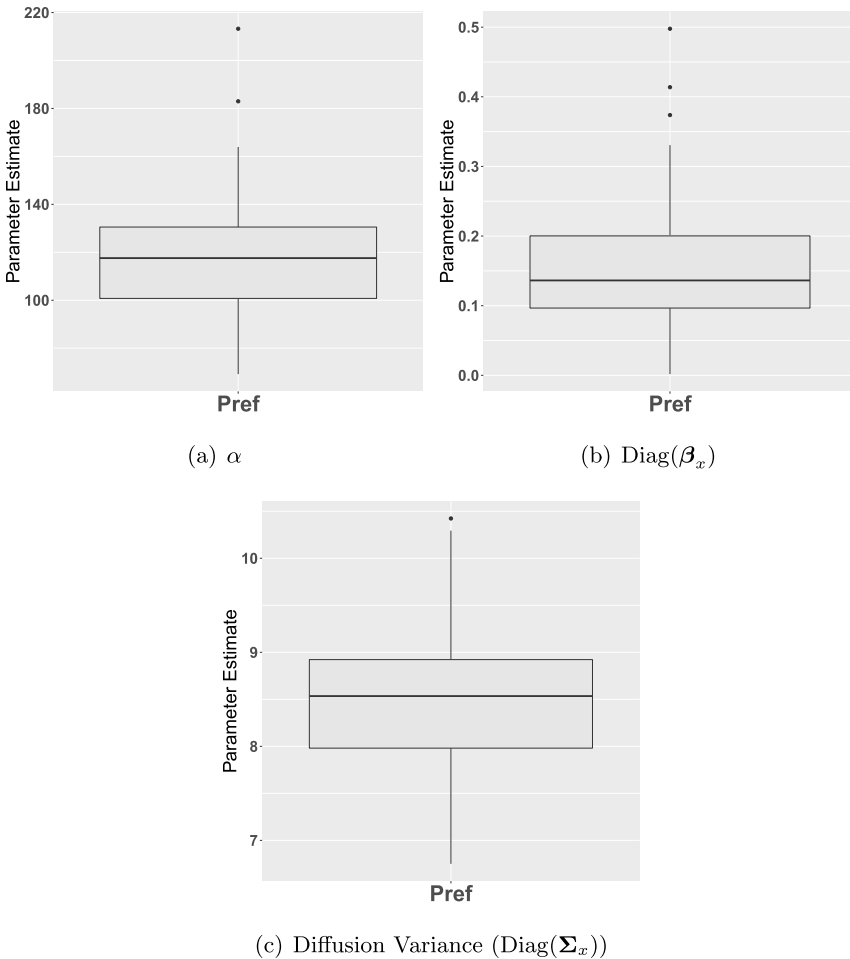


FIG. 9. Movement parameter estimates over 100 preferentially sampled simulated data sets.

generate the fine-time-scale data [Gurarie et al. (2017)], the boxplots in Figure 9 show that the estimators have a reasonable sampling distribution which seem to be unimodal and symmetric around their means.

D.3. Results for nonpreferential data. We also generated 100 data sets with nonpreferentially sampled data (i.e., setting $\alpha = 0$ in (3.7)), and estimated the parameters of the spatial process using the Preferential-CRW likelihood and the standard one (that conditions on the observed locations). The results in Figure 10 show

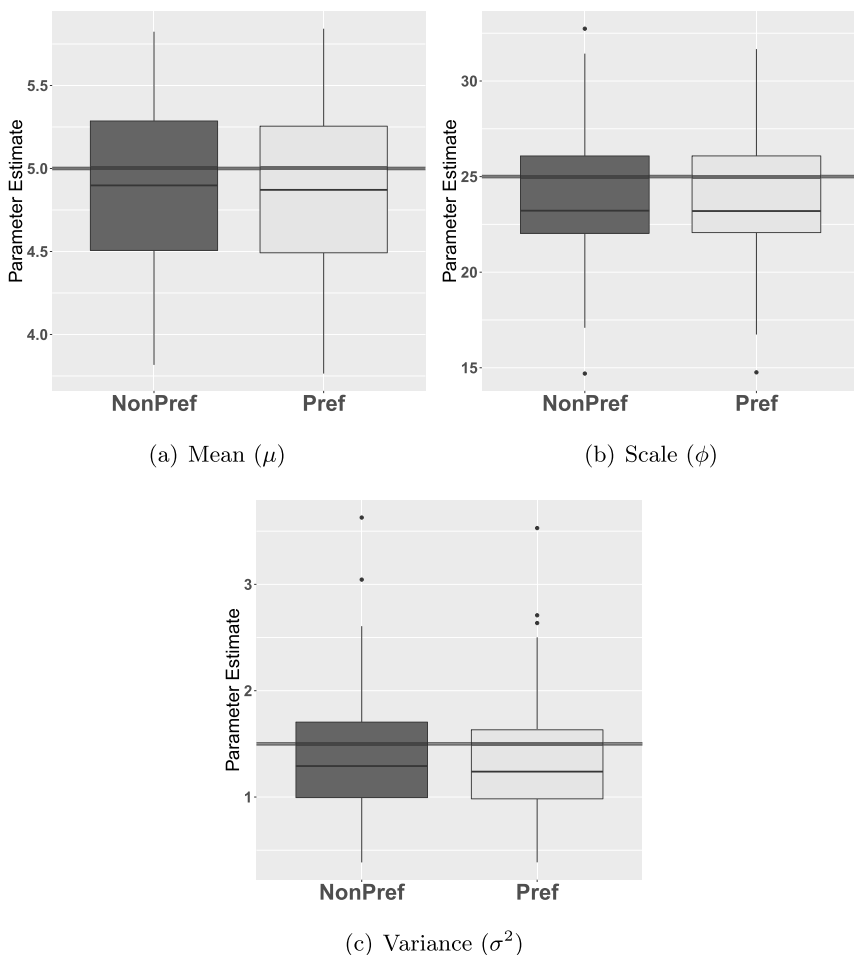


FIG. 10. Field parameter estimates over 100 nonpreferentially sampled simulated data sets with true parameter values marked as grey lines. The abbreviations NonPref and Pref stand for the standard MLE (nonpreferential) estimation and the one using the Preferential-CRW model of Section 3.1.

that, as expected, there is no practical difference between the parameter estimates obtained using either likelihood when no preferential sampling is present.

We also performed the same comparisons as above over the 100 nonpreferentially sampled data sets to verify that in this case, as expected, there was little qualitative difference between the predictions obtained using either method. The results are displayed in Figure 11 and show that, although RMSPE's are often larger with the preferential model, the difference between the two models is minor compared to the preferentially sampled data. This can be observed by comparing the scale of the differences in the plots, which are considerably smaller in the nonpreferentially sampled analysis.

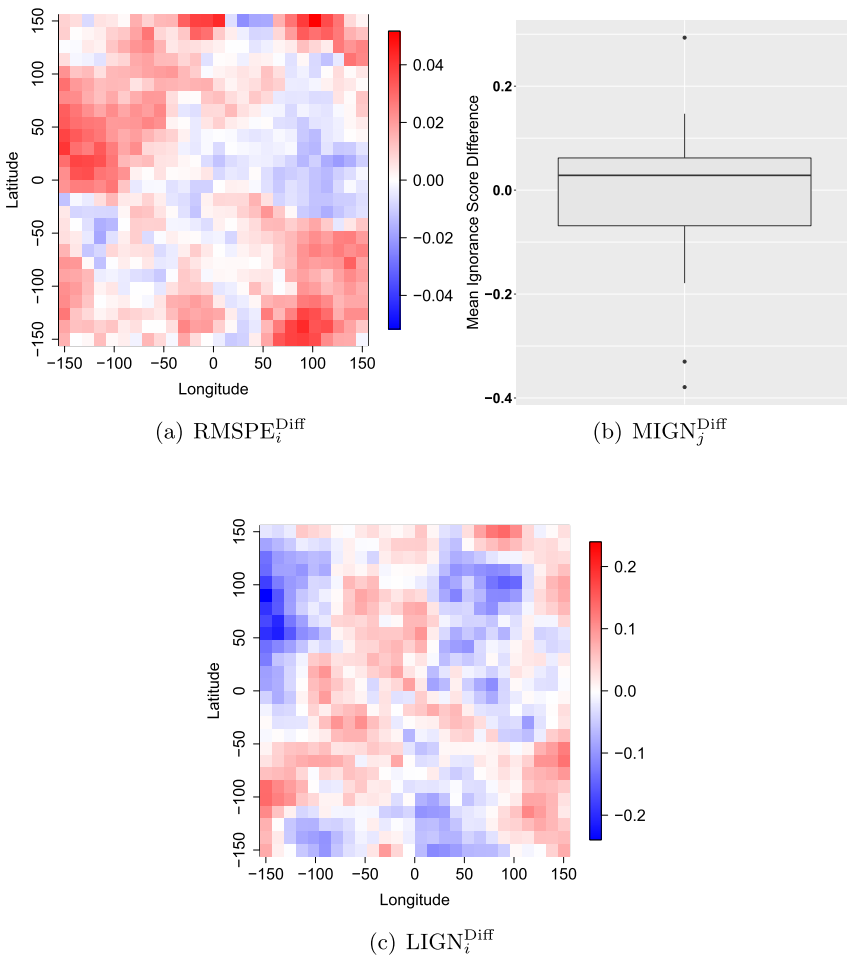


FIG. 11. Comparison of root mean square prediction error (RMSPE) difference, mean ignorance score (MIGN) difference and Location ignorance score (LIGN) difference respectively across 100 non-preferentially sampled simulated data sets.

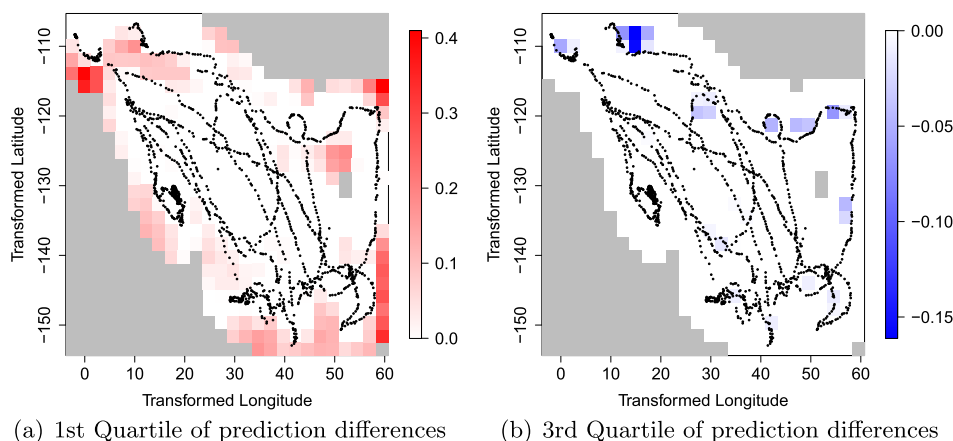


FIG. 12. *Quartiles of prediction differences over the 50 subsamples for the real data set analysis.*

APPENDIX E: MAGNITUDE OF PREDICTION DIFFERENCES

We have tried different ways to assess the significance of the prediction differences in Figure 7. Panel (b) shows the median difference, while the first (25%) and 3rd (75%) quartiles of the prediction differences (over the 50 subsamples) are shown in panels (a) and (b) of Figure 12. Note that in this plot, the white areas are locations of no-significance (negative-valued regions in the 25% quantile and positive regions in the 75% quantile). In other words, the coloured regions of the 25% quantile plot are regions in which prediction difference were consistently positive, and coloured regions in the 75% quantile plot consistently negative.

As it may be expected, in large parts of the domain the differences in predictions do not show a clear trend either way (white locations), but some conclusions could be drawn from these figures. The coloured regions in panel (a) are zones where most of the predictions for the preferential sampling model were higher than from the standard one. We see that they are in regions outside the area covered by the data, which is consistent with the positive estimates for α (see Figure 8 on page 733). Similarly, the coloured regions in panel (b) are zones where the predictions were mostly lower. Note that these areas are almost all among the observed locations, which is also consistent with our conclusions above.

SUPPLEMENTARY MATERIAL

Simulation Code (DOI: [10.1214/18-AOAS1217SUPP](https://doi.org/10.1214/18-AOAS1217SUPP); .zip). Code used for the simulations in this paper, with an example shown in the README.

REFERENCES

- ALBERTSEN, C. M., WHORISKEY, K., YURKOWSKI, D., NIELSEN, A. and FLEMMING, J. M. (2015). Fast fitting of non-Gaussian state-space models to animal movement data via template model builder. *Ecology* **96** 2598–2604.

- AUGER-MÉTHÉ, M., FIELD, C., ALBERTSEN, C. M., DEROCHER, A. E., LEWIS, M. A., JONSEN, I. D. and FLEMMING, J. M. (2016). State-space models' dirty little secrets: Even simple linear Gaussian models can have estimation problems. *Sci. Rep.* **6** 26677.
- AUGER-MÉTHÉ, M., ALBERTSEN, C. M., JONSEN, I. D., DEROCHER, A. E., LIDGARD, D. C., STUDHOLME, K. R., BOWEN, W. D., CROSSIN, G. T. and FLEMMING, J. M. (2017). Spatiotemporal modelling of marine movement data using Template Model Builder (TMB). *Mar. Ecol. Prog. Ser.* **565** 237–249.
- BANERJEE, S. and GELFAND, A. E. (2003). On smoothness properties of spatial processes. *J. Multivariate Anal.* **84** 85–100. [MR1965824](#)
- BANERJEE, S., GELFAND, A. E. and SIRMANS, C. F. (2003). Directional rates of change under spatial process models. *J. Amer. Statist. Assoc.* **98** 946–954. [MR2041483](#)
- BOLKER, B. M., GARDNER, B., MAUNDER, M., BERG, C. W., BROOKS, M., COMITA, L., CRONE, E., CUBAYNES, S., DAVIES, T., VALPINE, P. et al. (2013). Strategies for fitting nonlinear ecological models in R, AD model builder, and BUGS. *Methods Ecol. Evol.* **4** 501–512.
- BREED, G. A., JONSEN, I. D., MYERS, R. A., BOWEN, W. D. and LEONARD, M. L. (2009). Sex-specific, seasonal foraging tactics of adult grey seals (*Halichoerus grypus*) revealed by state-space analysis. *Ecology* **90** 3209–3221.
- BREED, G. A., COSTA, D. P., JONSEN, I. D., ROBINSON, P. W. and MILLS-FLEMMING, J. (2012). State-space methods for more completely capturing behavioral dynamics from animal tracks. *Ecol. Model.* **235–236** 49–58.
- BRILLINGER, D. R., STEWART, B. S. and LITTNAN, C. L. (2008). Three months journeying of a Hawaiian monk seal. In *Probability and Statistics: Essays in Honor of David A. Freedman. Inst. Math. Stat. (IMS) Collect.* **2** 246–264. IMS, Beachwood, OH. [MR2459954](#)
- BRILLINGER, D. R., PREISLER, H. K., AGER, A. A., KIE, J. G. and STEWART, B. S. (2002). Employing stochastic differential equations to model wildlife motion. *Bull. Braz. Math. Soc. (N.S.)* **33** 385–408. [MR1978835](#)
- BUDERMAN, F. E., HOOTEN, M. B., IVAN, J. S. and SHENK, T. M. (2016). A functional model for characterizing long-distance movement behaviour. *Methods Ecol. Evol.* **7** 264–273.
- CARTON, J. A. and GIESE, B. S. (2008). A reanalysis of ocean climate using simple ocean data assimilation (SODA). *Mon. Weather Rev.* **136** 2999–3017.
- CARTON, J. A., CHEPURIN, G. A., CHEN, L. and GRODSKY, S. A. (2018). Improved global net surface heat flux. *J. Geophys. Res.* **123** 3144–3163.
- CONN, P. B., THORSON, J. T. and JOHNSON, D. S. (2017). Confronting preferential sampling when analysing population distributions: Diagnosis and model-based triage. *Methods Ecol. Evol.* **8** 1535–1546.
- DA SILVA FERREIRA, G. and GAMERMAN, D. (2015). Optimal design in geostatistics under preferential sampling. *Bayesian Anal.* **10** 711–735. [MR3420820](#)
- DIGGLE, P. J., MENEZES, R. and SU, T. (2010). Geostatistical inference under preferential sampling. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **59** 191–232. [MR2744471](#)
- DIGGLE, P. J. and RIBEIRO, P. J. JR. (2007). *Model-Based Geostatistics. Springer Series in Statistics.* Springer, New York. [MR2293378](#)
- DINSDALE, D. and SALIBIAN-BARRERA, M. (2019a). Methods for preferential sampling in geostatistics. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **68** 181–198. [MR3902988](#)
- DINSDALE, D. and SALIBIAN-BARRERA, M. (2019b). Supplement to “Modelling ocean temperatures from bio-probes under preferential sampling.” DOI:[10.1214/18-AOAS1217SUPP](#).
- DUJON, A. M., LINDSTROM, R. T. and HAYS, G. C. (2014). The accuracy of fastloc-GPS locations and implications for animal tracking. *Methods Ecol. Evol.* **5** 1162–1169.
- EVANS, K., LEA, M.-A. and PATTERSON, T. (2013). Recent advances in bio-logging science: Technologies and methods for understanding animal behaviour and physiology and their environments. *Deep-Sea Res., Part 2, Top. Stud. Oceanogr.* **88** 1–6.

- FEDAK, M. (2004). Marine animals as platforms for oceanographic sampling: A situation for biology and operational oceanography. *Mem. Natl. Inst. Polar Res., Spec. Issue* **58** 133–147.
- FEDAK, M. A. (2013). The impact of animal platforms on polar ocean observation. *Deep-Sea Res., Part 2, Top. Stud. Oceanogr.* **88–89** 7–13.
- GELFAND, A. E., SAHU, S. K. and HOLLAND, D. M. (2012). On the effect of preferential sampling in spatial prediction. *Environmetrics* **23** 565–578. [MR3020075](#)
- GLOAGUEN, P., ETIENNE, M.-P. and LE CORFF, S. (2018). Stochastic differential equation based on a multimodal potential to model movement data in ecology. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **67** 599–619. [MR3787968](#)
- GNEITING, T. (2013). Strictly and non-strictly positive definite functions on spheres. *Bernoulli* **19** 1327–1349. [MR3102554](#)
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. [MR2345548](#)
- GOULD, J., ROEMMICH, D., WIJFFELS, S., FREELAND, H., IGNASZEWSKY, M., JIANPING, X., POULIQUEN, S., DESAUBIES, Y., SEND, U., RADHAKRISHNAN, K. et al. (2004). Argo profiling floats bring new era of in situ ocean observations. *Eos Trans. AGU* **85** 185–191.
- GRIEWANK, A. and WALTHER, A. (2008). *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, 2nd ed. SIAM, Philadelphia, PA. [MR2454953](#)
- GUINET, C., VACQUIÉ-GARCIA, J., PICARD, B., BESSIGNEUL, G., LEBRAS, Y., DRAGON, A. C., VIVIANI, M., ARNOULD, J. P. and BAILLEUL, F. (2014). Southern elephant seal foraging success in relation to temperature and light conditions: Insight into prey distribution. *Mar. Ecol. Prog. Ser.* **499** 285–301.
- GURARIE, E. and OVASKAINEN, O. (2011). Characteristic spatial and temporal scales unify models of animal movement. *Amer. Nat.* **178** 113–123.
- GURARIE, E., FLEMING, C. H., FAGAN, W. F., LAIDRE, K. L., HERNÁNDEZ-PLIEGO, J. and OVASKAINEN, O. (2017). Correlated velocity models as a fundamental unit of animal movement: Synthesis and applications. *Mov. Eco.* **5** 13.
- HOOTEN, M. B. and JOHNSON, D. S. (2017a). Basis function models for animal movement. *J. Amer. Statist. Assoc.* **112** 578–589. [MR3671753](#)
- HOOTEN, M. B., JOHNSON, D. S., MCCLINTOCK, B. T. and MORALES, J. M. (2017b). *Animal Movement: Statistical Models for Telemetry Data*. CRC Press, Boca Raton, FL.
- JACOBS, S. (2006). Observations of change in the Southern Ocean. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **364** 1657–1681.
- JEONG, J. and JUN, M. (2015). A class of Matérn-like covariance functions for smooth processes on a sphere. *Spat. Stat.* **11** 1–18. [MR3311853](#)
- JOHNSON, D. S., LONDON, J. M., LEA, M.-A. and DURBAN, J. W. (2008). Continuous-time correlated random walk model for animal telemetry data. *Ecology* **89** 1208–1215.
- JONSEN, I. D., FLEMMING, J. M. and MYERS, R. A. (2005). Robust state–space modeling of animal movement data. *Ecology* **86** 2874–2880.
- KRISTENSEN, K., NIELSEN, A., BERG, C. W., SKAUG, H. and BELL, B. M. (2016). TMB: Automatic differentiation and Laplace approximation. *J. Stat. Softw.* **70** 1–21. DOI:[10.18637/jss.v070.i05](#).
- LINDGREN, F. and RUE, H. (2015). Bayesian spatial modelling with R-INLA. *J. Stat. Softw.* **63** 1–25.
- 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. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **73** 423–498. [MR2853727](#)
- LIU, Y., BATAILLE, B. C., TRITES, A. W. and ZIDEK, J. V. (2015). Bias correction and uncertainty characterization of dead-reckoned paths of marine mammals. *Animal Biotelemetry* **3** 51–61.

- MCCLINTOCK, B. T., KING, R., THOMAS, L., MATTHIOPOULOS, J., MCCONNELL, B. J. and MORALES, J. M. (2012). A general discrete-time modeling framework for animal movement using multistate random walks. *Ecol. Monogr.* **82** 335–349.
- MCINTYRE, T., ANSORGE, I. J., BORNEMANN, H., PLÖTZ, J., TOSH, C. A. and BESTER, M. N. (2011). Elephant seal dive behaviour is influenced by ocean temperature: Implications for climate change impacts on an ocean predator. *Mar. Ecol. Prog. Ser.* **441** 257–272.
- MORALES, J. M., HAYDON, D. T., FRAIR, J., HOLSINGER, K. E. and FRYXELL, J. M. (2004). Extracting more out of relocation data: Building movement models as mixtures of random walks. *Ecology* **85** 2436–2445.
- PATI, D., REICH, B. J. and DUNSON, D. B. (2011). Bayesian geostatistical modelling with informative sampling locations. *Biometrika* **98** 35–48. [MR2804208](#)
- PREISLER, H. K., AGER, A. A. and WISDOM, M. J. (2013). Analyzing animal movement patterns using potential functions. *Ecosphere* **4** 1–13.
- ROBUSTO, C. C. (1957). Classroom Notes: The Cosine-Haversine Formula. *Amer. Math. Monthly* **64** 38–40. [MR1529506](#)
- ROQUET, F., WUNSCH, C., FORGET, G., HEIMBACH, P., GUINET, C., REVERDIN, G., CHARRASSIN, J.-B., BAILLEUL, F., COSTA, D. P., HUCKSTADT, L. A. et al. (2013). Estimates of the Southern Ocean general circulation improved by animal-borne instruments. *Geophys. Res. Lett.* **40** 6176–6180.
- ROQUET, F., WILLIAMS, G., HINDELL, M. A., HARCOURT, R., MCMAHON, C., GUINET, C., CHARRASSIN, J.-B., REVERDIN, G., BOEHME, L., LOVELL, P. et al. (2014). A southern Indian Ocean database of hydrographic profiles obtained with instrumented elephant seals. *Sci. Data* **1** 140028.
- ROULSTON, M. S. and SMITH, L. A. (2002). Evaluating probabilistic forecasts using information theory. *Mon. Weather Rev.* **130** 1653–1660.
- RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 319–392. [MR2649602](#)
- RUSSELL, J. C., HANKS, E. M., HARAN, M. and HUGHES, D. (2018). A spatially varying stochastic differential equation model for animal movement. *Ann. Appl. Stat.* **12** 1312–1331. [MR3834305](#)
- SHADDICK, G. and ZIDEK, J. V. (2014). A case study in preferential sampling: Long term monitoring of air pollution in the UK. *Spat. Stat.* **9** 51–65. [MR3326831](#)
- SIEGERT, S., FERRO, C. A. and STEPHENSON, D. B. (2014). Evaluating ensemble forecasts by the ignorance score—correcting the finite-ensemble bias. Preprint. Available at [arXiv:1410.8249](#).
- 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. [MR2873784](#)
- UNGAR, E. D., HENKIN, Z., GUTMAN, M., DOLEV, A., GENIZI, A. and GANSKOPP, D. (2005). Inference of animal activity from GPS collar data on free-ranging cattle. *Rangeland Ecology and Management* **58** 256–266.
- VOTIER, S. C., BEARHOP, S., WITT, M. J., INGER, R., THOMPSON, D. and NEWTON, J. (2010). Individual responses of seabirds to commercial fisheries revealed using GPS tracking, stable isotopes and vessel monitoring systems. *J. Appl. Ecol.* **47** 487–497.
- WEIMERSKIRCH, H., BONADONNA, F., BAILLEUL, F., MABILLE, G., DELL’OMO, G. and LIPP, H.-P. (2002). GPS tracking of foraging albatrosses. *Science* **295** 1259.
- WHORISKEY, K., AUGER-MÉTHÉ, M., ALBERTSEN, C. M., WHORISKEY, F. G., BINDER, T. R., KRUEGER, C. C. and FLEMMING, J. M. (2017). A hidden Markov movement model for rapidly identifying behavioral states from animal tracks. *Ecol. Evol.* **7** 2112–2121.
- WILSON, R. and WILSON, M.-P. (1988). Dead reckoning: A new technique for determining penguin movements at sea. *Meeresforschung* **32** 155–158.

DEPARTMENT OF STATISTICS
FACULTY OF SCIENCE
UNIVERSITY OF BRITISH COLUMBIA
3182 EARTH SCIENCES BUILDING, 2207 MAIN MALL
VANCOUVER, BRITISH COLUMBIA
CANADA V6T 1Z4
E-MAIL: danielrdinsdale@gmail.com
matias@stat.ubc.ca