



Estimating space-use and habitat preference from wildlife telemetry data

Geert Aarts, Monique MacKenzie, Bernie McConnell, Mike Fedak and Jason Matthiopoulos

G. Aarts, B. McConnell and M. Fedak, NERC Sea Mammal Research Unit, Gatty Marine Laboratory, Univ. of St. Andrews, St. Andrews, Fife KY16 8LB, Scotland, UK. – M. Mackenzie, School of Mathematics and Statistics, Univ. of St. Andrews, St. Andrews, Fife KY16 9SS, Scotland, UK, and Centre for Research into Ecological and Environmental Modelling, The Observatory, Univ. of St. Andrews, St. Andrews, Fife KY16 9LZ, Scotland, UK. – J. Matthiopoulos (jm37@st-andrews.ac.uk), NERC Sea Mammal Research Unit, Gatty Marine Laboratory, Univ. of St. Andrews, St. Andrews, Fife KY16 8LB, Scotland, UK, and Centre for Research into Ecological and Environmental Modelling, The Observatory, Univ. of St. Andrews, St. Andrews, Fife KY16 9LZ, Scotland, UK.

Management and conservation of populations of animals requires information on where they are, why they are there, and where else they could be. These objectives are typically approached by collecting data on the animals' use of space, relating these positional data to prevailing environmental conditions and employing the resulting statistical models to predict usage at other geographical regions. Technical advances in wildlife telemetry have accomplished manifold increases in the amount and quality of available data, creating the need for a statistical framework that can use them to make population-level inferences for habitat preference and space-use. This has been slow-in-coming because wildlife telemetry data are spatio-temporally autocorrelated, often unbalanced, presence-only observations of behaviourally complex animals, responding to a multitude of cross-correlated environmental variables.

We review the evolution of regression models for the analysis of space-use and habitat preference and outline the essential features of a framework that emerges naturally from these foundations. This allows us to derive a relationship between usage of points in geographical space and preference of habitats in environmental space. Within this framework, we discuss eight challenges, inherent in the spatial analysis of telemetry data and, for each, we propose solutions that can work in tandem. Specifically, we propose a logistic, mixed-effects approach that uses generalized additive transformations of the environmental covariates and is fitted to a response data-set comprising the telemetry and simulated observations, under a case-control design.

We apply this framework to a non-trivial case-study using satellite-tagged grey seals *Halichoerus grypus* from the east coast of Scotland. We perform model selection by cross-validation and confront our final model's predictions with telemetry data from the same, as well as different, geographical regions. We conclude that, despite the complex behaviour of the study species, flexible empirical models can capture the environmental relationships that shape population distributions.

Animals meet their requirements for survival, growth and reproduction by exploiting available resources within the restrictions imposed by their physiology and environment. Management and conservation of animal populations require a scientific understanding of this process attained only by viewing it in its spatial context (Levin 1992). That is why an increasing number of papers in applied ecology are concerned with where a particular organism is (Kernohan et al. 1998, Blundell et al. 2001, Matthiopoulos 2003a, Matthiopoulos et al. 2004), why it is there (Johnson 1980, Arthur et al. 1996, Boyce and McDonald 1999, Manly et al. 2002, Calenge et al. 2005) and where else it is likely to be (Buckland and Elston 1993, Guisan and Zimmermann 2000, Guisan et al. 2002, Scott et al. 2002). These objectives are typically approached by collecting data on the animals' use of space, relating these observations to

prevailing environmental conditions and employing these relations to predict usage at other points in space.

Often, space-use data are obtained from transect surveys (Buckland et al. 1993). These record animals in the vicinity of a set of sampling lines or points and therefore tend to yield relatively few sightings, particularly for rare species living in inaccessible environments. Alternatively, wildlife tracking techniques focus on the individual animal. By making full use of recent advances in communication and information technology radio- and satellite-telemetry have dramatically increased the volume of data on animal usage (White and Garrott 1990, Priede and Swift 1992, Fedak et al. 2002).

The concept of habitat preference attempts to quantify the inherent needs of animals, as expressed in the environment in which they were observed (Johnson 1980, Manly

et al. 2002). Studies on preference originally used hypothesis testing to compare sets of contrasting environmental conditions. More recently, environmental variables have been incorporated as covariates in spatial models of usage (Buckland and Elston 1993, Boyce and McDonald 1999, Nielsen et al. 2005). Because of their perceived potential for prediction, these statistical developments are fuelling the explosive increase of quantitative analyses in applied, spatial ecology (Guisan and Zimmermann 2000, Calenge 2006, Latimer et al. 2006).

However, these analyses are faced with inherent problems such as the cross-correlation between environmental variables, spatial autocorrelation in animal distribution, variable detectability of animals in different environments, various imbalances in sampling effort, unequal accessibility of different points in space and the animals' complex responses to their environment. Furthermore, the increasing power of telemetry methods presents additional demands for covariate data at the appropriate temporal and spatial scale and the focus of telemetry studies on a few individuals presents non-trivial challenges for population-level inference. Although these problems are shared by most studies of space-use and habitat preference there have been few comprehensive attempts to identify them, review existing solutions and draw on recent advances in statistical modelling for new ones.

This study firstly aims to outline each of these problems, their consequences and a combination of solutions within a single statistical framework. Secondly, it aims to illustrate this framework using an extensive set of satellite telemetry data collected from a sub-population of grey seals *Halichoerus grypus* off the east coast of the United Kingdom.

We draw heavily from three areas of research: general spatial modelling, the study of habitat preference and the analysis of individual movement, each of which has, rather confusingly, developed its own terminology (Hall et al. 1997). To aid clarity and consistency, at least within this paper, we precede the main body of the text with a section on definitions. We then review the evolution of regression models for the analysis of usage and preference and offer some new ideas that lay the foundations for our methodological approach. In the larger two sections that follow, we correspondingly address the paper's two main objectives. Finally, we discuss issues of spatial prediction and place our work in a wider ecological and methodological context.

Definitions

Analyses of geographical space consider the dimensions of latitude and longitude, and sometimes also elevation or depth. The spatial distribution of a species is the density of animals (or, their usage) over geographical space. Environmental space comprises multiple dimensions representing different biotic and abiotic environmental variables, or candidate covariates. We define an environment as a combination of conditions – a single point in environmental space. Although points in geographical space always have a counterpart in environmental space, the reverse is not true because several combinations of environmental conditions simply may not exist in nature.

The multi-dimensional envelope in environmental space in which an animal does or may live, is the realized or fundamental niche, respectively (Hutchinson 1957). The term habitat is more ambiguous. Traditionally (Whittaker and Root 1973) and in common use, it signified the entire region in geographical or environmental space occupied by a particular organism. This species-specific definition is not very useful for modelling usage or preference because it does not account for gradations in usage within a habitat. Alternatively, habitat can be defined as any collection of environments or a cuboid in environmental space. We prefer this latter definition of habitat because it allows comparisons between habitats, both within and across species. Under the coarsest classification of habitats, all environments belong to the same habitat. In the finest classification scheme, each environment is a unique habitat and so is every point in geographical space.

Johnson (1980) defines selection as “the process in which an animal actually chooses a resource or habitat”. Attempts to quantify selection as a function of environmental variables have led to the development of resource selection functions (RSFs-reviewed in Boyce and McDonald 1999, and Manly et al. 2002). A resource can be permanently or temporarily depleted by the activity of an organism. RSFs relate a species' spatial distribution to environmental variables that may or may not be depleted by the animals, so “habitat selection function” might have been a more intuitive name. The response variable in RSFs is said to be proportional to the probability of use but, in practice, RSFs are used to model the disproportionality between usage and availability, and not usage itself.

In contrast, early work defined preference as the likelihood of a resource or habitat being chosen if offered to the animals in equal amounts (Johnson 1980). In wild animals, which experience unequal habitat availability, preference is defined as the disproportionality between usage and availability (used in RSF, Manly et al. 2002).

With all of the above in mind, we adopt the following working definition: given a habitat classification scheme (including the two trivial schemes outlined above), habitat preference is the ratio of the use of a habitat over its availability, conditional on the availability of all habitats to the study animals.

The statistical analysis of habitat preference

Hypothesis testing

Given sufficient time for movement, the expected spatial distribution of a population of unconstrained random walkers within a region of space is approximately uniform. By definition, random walkers show no habitat preference. Their uniform spatial distribution means that they use each habitat approximately in proportion to its availability, the total area occupied by the habitat in geographical space. Deviations from direct proportionality between usage and availability are interpreted as the tell-tale sign of habitat preference. The null model of proportional use can yield selectivity indexes for the i th habitat such as the ratio (w_i)

$$w_i = \frac{u_i}{a_i} \quad (a_i > 0) \quad (1)$$

of the expected number (u_i) of wildlife telemetry observations occurring in the i th habitat over the total area (a_i) taken up by that habitat (Manly et al. 2002). Note that this index of preference includes both habitats that are avoided and preferred. This can be analyzed by means of chi-squared tests (Allredge and Ratti 1992) or classic parametric techniques such as multivariate analysis of variance (MANOVA) (Aebischer et al. 1993).

Regression with discrete covariates

The MANOVA framework can be interpreted as a special case of regression with discrete covariates (Agresti 1996, Fox 1997). The number of telemetry locations observed in a particular habitat (\hat{u}_i) is a count and, if for the purposes of data collection, storage, or analysis, geographical space is represented by a grid, then the number of grid-cells belonging to a particular habitat is also a count. Therefore, the index of preference in eq. (1) will always be non-negative and usually a rational number and will have a skewed distribution for low counts. Therefore, the number of telemetry observations in the i th habitat can be modelled as a heterogeneous Poisson process (Manly et al. 2002) with rate $u_i = a_i w_i$

$$\begin{aligned} \hat{u}_i &\sim \text{Poisson}(a_i w_i) \\ a_i w_i &= a_i g^{-1}(\eta_i) = a_i e^{\eta_i} = e^{\log(a_i) + \eta_i}, \\ \eta_i &= \beta_0 + \beta_1 x_{i,1} + \dots + \beta_j x_{i,j} \end{aligned} \quad (2)$$

where η_i is the linear predictor, $g^{-1}(\cdot)$ is the inverse link function, x denotes a discrete or discretized environmental variable and x_{ij} denotes the value taken by the j th environmental variable in the i th habitat under the habitat classification scheme employed. $\log(a_i)$ is known as the model's offset. The fact that the coefficients β_j are subscripted by environmental variable and not by habitat, hints at why eq. (2) is a predictive model: it describes a trend across environmental space and provides estimates for the response variable in unobserved habitats. Equation (2) is a generalized linear model (GLM) whose parameters can be estimated by likelihood methodology (Agresti 1996, Fox 1997).

This approach still requires a habitat classification. For some environmental variables (e.g. vegetation type) this classification may be inherent in the way the data are collected, but for variables measured on a continuous scale (e.g. temperature, depth), discrete classes (or bins) are often constructed arbitrarily, prior to modelling. This can severely prejudice the results of the analysis: making the classification too coarse merges habitats that appear similar to a human observer but may be perceived differently by the animals. In contrast, too fine a classification, may lead to a plethora of habitats that are either non-existent in geographical space or contain no animal usage. This becomes more acute in studies with many environmental variables because the data become more thinly spread over a higher-dimensional environmental space. Possible remedies include using a zero-inflated or otherwise over-dispersed Poisson error (Welsh et al. 1996, Fox 1997) or negative binomial distributions (Nielsen et al. 2005), or by abandoning

classification in geographical and environmental space (see below).

Generalized linear models with continuous covariates

Our Supplementary material demonstrates how the discretization of geographical and environmental space can be replaced by a case-control design, similar to logistic-regression methods used elsewhere (Boyce and McDonald 1999). However, the method proposed here, links control observations with tagged individuals and selects the values of the environmental conditions at times corresponding to the actual telemetry observations. The case-control design produces a binomial response variable (\hat{u}_i) which takes the values 1 for i th ($i = 1, 2, \dots, n$) data point if it belongs to the telemetry data, and 0 for a control data point. This can be modelled as a Bernoulli process with probability h_i

$$\begin{aligned} \hat{u}_i &\sim B(1, h_i) \\ h_i &= g^{-1}(\eta_i) = \frac{e^{\eta_i}}{1 + e^{\eta_i}} \\ \eta_i &= \beta_0 + \beta_1 x_{i,1} + \dots + \beta_j x_{i,j} \end{aligned} \quad (3)$$

There are two important points about the response variable in eq. (3) when fitted to case-control data. First, because the number of absences used to fit the model is determined arbitrarily, the case-control approach quantifies the relative importance of different covariates but not the absolute abundance of animals. Second, Keating and Cherry (2004) have recently argued against the suggestion (Boyce and McDonald 1999, Manly et al. 2002) that, in logistic models, h_i is proportional to expected space-use. We argue (Supplementary material) that, for a large number of controls, h_i is, in fact, proportional to preference.

To derive usage from h_i we follow an approach similar to that of Lele and Keim (2006) and Johnson et al. (2006). We assume that the telemetry observations ($\hat{u}_i = 1$) are generated from a heterogeneous, spatial Poisson process whose rate is proportional to the unknown, spatial probability density function $f_u(X_s)$, where X_s is a vector of values for the environmental variables at position s in geographical space (Fig. 1a). Similarly, the control observations ($\hat{u}_i = 0$) are generated from a user-defined spatial Poisson process with a rate $f_a(X_s)$ (Fig. 1b). In the simplest case, where all points within a study area are assumed to be equally accessible, $f_a(X_s) = A^{-1}$, where A is the total area of the study region. More complex spatial density functions can be used when such an approximation is inappropriate, e.g. for central-place foragers (see also Problem 2 below).

By indefinitely increasing the resolution of both the spatial and environmental grids (operations O1 and O2 in Supplementary material), the expected number of telemetry observations in the i th habitat u_i can be approximated by $k_u f_u(X_s) c$, where c is the area of a geographical grid cell and the proportionality constant k_u can be thought of as the expected size of the telemetry data set and depends on the number of animals tagged and frequency of telemetry locations. Similarly, the expected number of controls in the i th habitat $k_a p_i$ (Supplementary material), can be approximated by $k_a f_a(X_s) c$, where k_a is the user-defined quantity of the total number of absences




Fig. 1. The paper’s fundamental concepts illustrated with an artificial example. The underlying, unknown surface of space-use $f_u(X_s)$, in (a), generated the telemetry observations ($\hat{u} = 1$), shown as dots. Pseudo-absence data ($\hat{u} = 0$) in (b) were generated from a spatial rate $f_a(X_s)$, representing the accessibility from the animals’ nominal central place. The distribution of usage in (a) is caused by a single, heterogeneously distributed environmental covariate X_s (c). This relationship is modelled by the response h in one-dimensional environmental space (d). Predictions of the response variable in geographical space (e) give a visual representation of preference. The simulated data in this example, allowed us to perform a spatial comparison between predicted usage $f_u(X_s)$ (filled contours) and true, unobserved usage $\tilde{f}_u(X_s)$ (bold contour lines) (f).

generated in space. Following O1 and O2 the response variable $h(X_s)$ in eq. (3) tends to

$$h(X_s) = \frac{k_v f_u(X_s)}{k_a f_a(X_s) + k_v f_u(X_s)}. \quad (4)$$

This probability $h(X_s)$ depends on the values for the environmental covariates $X_s = (x_1, \dots, x_j, \dots)$ observed at spatial location \mathbf{s} (Fig. 1c). These environmental conditions vary spatially (Fig. 1d) and so does $h(X_s)$ (Fig. 1e). Finally, we can use the dependence between h and X_s , to estimate spatial usage at location \mathbf{s} . Equation (4) can be re-arranged as

$$f_u(X_s) = \frac{h(X_s)}{1 - h(X_s)} r f_a(X_s), \quad (5)$$

where, $r = k_a/k_u$ is the number of controls selected per telemetry observation. This is a local estimate of usage. Since it is only based on a random sample of controls, the cumulative estimated usage over all space will usually deviate from unity and will need to be normalized. Experimentation with test data sets has indicated that this does not bias the estimates of usage. This would still need to be proven theoretically as part of future work. The normalized map of usage can be provide the proportion of animals within a specified region (e.g. a special area of conservation (SAC)), by integrating $f_u(X_s)$ over that region. If the total number of animals in the region is also known from other sources, relative usage can be scaled up into a map of absolute population density.

Preference is then seen as an exponential transformation of the linear predictor (Lele and Keim 2006, Johnson et al. 2006), multiplied by a constant which accounts for the unequal number of cases and controls selected

$$w(X_s) = \frac{f_u(X_s)}{f_a(X_s)} = \frac{h(X_s)}{1 - h(X_s)} r = \frac{e^{\eta_i}}{1 + e^{\eta_i}} r = \frac{e^{\eta_i} r}{1 + e^{\eta_i}} = e^{\eta_i + \ln(r)}. \quad (6)$$

The case-control design raises the possibility of false absences (Boyce et al. 2002), a serious problem for small-scale studies on coarse spatial grids. However, the large spatial scales of telemetry data, allow us to treat individuals as practically dimensionless and the probability of encountering a tagged animal at the exact coordinates of a given geographic location as practically zero.

Modelling habitat preference: problems and solutions

The quality of a statistical model is determined by its accuracy, precision and parsimony (Burnham and Anderson 2002). Specific properties of the environment, the characteristics of the study animals and the ways that both are sampled, can cause one or more of these characteristics to suffer (Levins 1966). Loss of accuracy implies biased parameter estimates and predictions, and loss of precision leads to increased uncertainty in parameters and predictions. Loss in parsimony leads to over-parameterized models capable of predicting a particular data set well, but liable to predict new data poorly.

Below, we discuss eight problems that can cause such detrimental effects and propose appropriate solutions.

Problem 1: environmental data rarely coincide with usage data

In relating wildlife telemetry data to environmental variables it is implicitly assumed that the study animals experience the same conditions described by our environmental data. However, environmental data are usually collected independently of wildlife telemetry data and unwanted interferences (e.g. cloud cover obscuring remote sensing) or logistical constraints (e.g. limited observation time) can lead to incomplete spatial coverage. Consequently, it is rarely known exactly what conditions the animals are responding to, at any given point in space and time. This reduces model precision, increasingly so with highly heterogeneous and dynamic environmental variables (Isaaks and Srivastava 1990).

To address the lack of environmental data some modern tags also collect data about the animals' environment (Fedak 2004), but such valuable technological improvements tell us nothing about conditions at points not visited by the tagged animals. This information is just as important for the analysis of preference.

If the errors produced by the environmental survey method are negligible then an interpolated surface that is constrained to pass through the survey observations, can be used (Ripley 1981). However, if, the environmental data contain measurement error (e.g. data on prey density) or process stochasticity (e.g. meteorological variables), smoothing techniques are more appropriate estimation methods (Ripley 1981, Silverman 1986). Approaches such as kernel-smoothing (Silverman 1986), attempt to reproduce the mean and underlying distribution of the stationary process that generated the survey data. Interpolation and smoothing can be seen as the two extremes of the methodological spectrum for density estimation. A third estimation technique, kriging (Isaaks and Srivastava 1990), first models spatial autocorrelation in the form of the variogram which it then uses for estimation. An appealing aspect of kriging is that it can behave as a spatial interpolator as well as a smoother depending on the variogram intercept.

Problem 2: points in space are not equally accessible to the animals

The precise definition of habitat availability is important when modelling preference (eq. 1). Measuring it as the total area taken up by a particular habitat either implies that animals have access to all points in geographical space equally (perfect mixing) or that, within their sensory range, they always experience a representative sample of habitats (representative perception) (Garshelis 2000).

Often studies of preference take a pragmatic approach, focusing on arbitrarily-defined regions of geographical space (Manly et al. 2002). Others (Johnson 1980) identified accessible space as the estimated animal's home range. This still assumes that all points within the home range are equally accessible from its centre and that accessible points outside the observed home range are simply not preferred.

Other researchers have taken a more mechanistic approach by calculating the accessibility of points in space from the animals' starting position, speed and mode of movement, travel duration and travel medium/obstacles. Such models have been developed for both nomadic (Arthur et al. 1996, Hjermmann 2000) and central-place foragers (Matthiopoulos 2003b), and can be parameterized from readily available data. Their output is a spatial surface, representing the likelihood of observing an animal at a given point, given no habitat preference. It can be treated as a probability density function $f_a(\cdot)$ from which random points can be sampled to construct the set of absence points for eq. (3). Since preference is defined as the likelihood of a habitat being chosen if offered to the animal in equal amount, selecting the absence data from an accurate accessibility surface keeps the response variable (eq. (3) and (4)) in agreement with the definition of preference and eq. (1).

Certainly, accessibility surfaces will never be perfect and discrepancies between the response variable and true preference will cause the model to over- or under-predict. To absorb these residuals some measure of accessibility (e.g. distance from the central place) can be incorporated into the model as a candidate covariate.

Problem 3: sampling effort in telemetry studies is rarely balanced across individuals

With the exception of studies of rare species, which may be interested in the particular animals carrying the telemetry tags, most habitat preference studies aim to make inferences about the behaviour of the entire population to which the tagged animals belong. All population-level inferences are subject to sampling error. In telemetry studies, sampling error is usually large because, due to logistical constraints, only a small sample of animals are tagged and because sampling effort between tagged individuals is almost never balanced. Also, capture and tagging effort may not be spatially uniform and might not sample animals of different ages and sex representatively. Therefore, estimating habitat preference by pooling telemetry data from all individuals is likely to bias the results towards certain data-rich individuals, types of individuals or regions of geographical space. The alternative, is to recognize the natural hierarchy of sampling units (Fig. 2) and use an error structure that more accurately reflects the variability within and between different levels of this hierarchy (Thomas and Taylor 2006, Gillies et al. 2006). Multi-level or mixed-effects models (Pinheiro and Bates 2000, Fox 2002), can simultaneously model the behaviour of the average individual using the population mean (fixed-effect) and the variability in the behaviour across individuals using random effects. Equation (3) can be modified into a mixed-effects model as follows:

$$\hat{u}_i \sim B(1, h_i)$$

$$h_i = g^{-1}(\eta_i) = \frac{e^{\eta_i}}{1 + e^{\eta_i}} \quad (7)$$

$$\eta_i = b_{0,m} + b_{1,m}x_{i,1} + \dots + b_{j,m}x_{i,j} + \dots$$

The coefficients b are, themselves, random variables that can either be specified with a mean and variance, or modelled as functions of class-member characteristics. For




Fig. 2. Illustration of the hierarchy of sampling units in telemetry data. The population (a) comprises sub-populations (b), each of which contains several individuals (c). Every individual makes trips (d) to sea during which telemetry observations (e) are collected.

example, to capture the individual variation within a group of animals, the class must be defined as the group of individuals. Class-member characteristics appropriate for that class might be an individual's age, sex or mass. Hence, the coefficient $b_{j,m}$ that quantifies how the m th individual responds to the j th environmental variable can be given as a linear combination of individual-specific characteristics ($z_{k,m}$) using coefficients (β_j) that refer to the entire group

$$b_{j,m} = \beta_{j,0} + \beta_{j,1}z_{1,m} + \cdots + \beta_{j,k}z_{k,m} + \cdots v_j. \quad (8)$$

Individual characteristics ($z_{k,m}$) enter the linear predictor as interactions with the environmental covariate ($x_{i,j}$)

$$\eta_i = b_{0,m} + b_{1,m}x_{i,1} + \cdots + (\beta_{j,0}x_{i,j} + \beta_{j,1}z_{1,m}x_{i,j} + \cdots + \beta_{j,k}z_{k,m}x_{i,j} + \cdots + vx_{i,j}) \cdots. \quad (9)$$

The coefficients β are also known as the fixed effects. The random effect accounts for the within-class (in our application, between-individual) variability. They are denoted as $v_0, v_1, \dots, v_j, \dots$ and are commonly assumed to have a joint multivariate normal distribution with mean zero and a variance-covariance matrix Ψ , representing within-class variability (Pinheiro and Bates 2000). The estimation procedure for mixed-effects model returns values for the fixed effects and estimates of Ψ for the distribution of the random effects.

In some cases, the responses of individuals to environmental variables are distributed non-normally around the population mean response. Specifically, a few individuals may have extreme responses (implying a non-normal kurtosis for the random effect) or their responses may be asymmetric around the average population response (giving rise to a skewed distribution for the random effect). These two deviations from normality cause imprecision in the estimates of variance for the random effects and biases in the estimates of the fixed effects, respectively. If such deviations are a consistent result of identifiable characteristics of the individual then these can be included in the model to explicitly account for individual variation and yield normally distributed random effects. If they are unknown, a non-normal random-effect distribution may be required.

Problem 4: some environmental variables may be cross-correlated

Certain large-scale processes (e.g. meteorological or geological) may influence most of the environmental variables that might be used to explain usage. Furthermore, interactions between environmental variables are often just as strong as the links between them and usage. Both of these mechanisms may lead to correlations between one candidate covariate and another (colinearity) or a linear combination of others (multi-colinearity) (Cramer 1985).

A corollary of the existence of multi-colinearity in multi-variable models is that the dimensionality of the model is not entirely supported by the data. Using such an unparsimonious model results in unstable parameter estimates that may have large standard errors and be sensitive to outliers (Fox 1997). Colinearity needs to be detected and treated, particularly if biological interpretation of the model's parameter estimates and associated variances are

important and if the model is to be used for predictions elsewhere.

The simpler cross-correlations are usually detected by inspecting pair-wise plots. Multi-colinearities can be detected with variance inflation factors (VIFs) (Fox 1997). Recent approaches, using GLMs, have also treated the consequences of multi-colinearity by using VIFs to adjust the standard errors of model parameters (Fox 1997). This is a promising approach and we look forward to its implementation for mixed-effect models.

Other treatments of colinearity involve obtaining uncorrelated transformations of the candidate covariates by means of techniques such as principal components analysis (Jolliffe 1990). A disadvantage of this approach is that the principal components of several environmental variables are difficult to interpret biologically. Alternatively, colinearity can be treated by dropping as many environmental variables as required, by means of model-selection criteria such as the change in deviance, approximate F-tests, or information criteria (ICs) (McCullagh and Nelder 1989, Hastie and Tibshirani 1990, Burnham and Anderson 2002). Automatic model-selection may be supervised by using existing biological knowledge on the causal relationships between variables.

Problem 5: species distributions are spatially autocorrelated

Positive spatial autocorrelation is a typical characteristic of animal distributions. It leads to nearby geographic locations having more similar values of usage than would be expected by chance. In telemetry data, this is manifested as clusters of observations in space which, in itself, is not problematic. In fact, this interdependence between points in space is usefully employed by all usage estimation techniques (such as interpolation, smoothing and kriging, discussed above) (Blundell et al. 2001, Matthiopoulos 2003a). In habitat preference studies, it is hoped that autocorrelation in usage is a result of autocorrelation in the available covariates and therefore that it will be captured by the model. However, sometimes the best-fitting model presents residual autocorrelation: it systematically over/under-estimates usage in entire regions of geographical space. Such residual (positive) autocorrelation violates the central assumption of independence leading to under-estimates of the standard errors for the parameters (a false increase in model precision) and a more likely inclusion of irrelevant environmental variables (i.e. a loss in model parsimony).

Residual autocorrelation can be detected by Moran's I, Geary's C (Cliff and Ord 1973), or by constructing spatial variograms of the model residuals (Isaaks and Srivastava 1990). If there is no significant residual spatial autocorrelation, no action needs to be taken. If residual spatial autocorrelation is detected, it could be due to either extrinsic (e.g. autocorrelated environment) or intrinsic (e.g. conspecific attraction, dispersal limitations) factors (Legendre 1993, Keitt et al. 2002, Overmars et al. 2003).

Extrinsic factors give rise to residual autocorrelation either because an important, autocorrelated covariate has been omitted from the model, or due to model mis-specification

(Cliff and Ord 1973). So, when extrinsic causes are suspected, the first step is to introduce new covariates and ensure that the model is sufficiently flexible. Persistent residual autocorrelation may warrant the use of conditional and simultaneous autoregressive models (CAR and SAR, respectively – Keitt et al. 2002), also known as models with spatially filtered variables, or geostatistical models such as co-kriging (Stein and Corsten 1991).

An intrinsic cause means that observing an animal at a point in space is a direct consequence of the presence of conspecifics at neighbouring points. Alternatively, spatial auto-correlation may be the result of temporal auto-correlation within individual data sets from single individuals (see below). Remaining intrinsic auto-correlation (e.g. due to presence of conspecifics) can be dealt with by augmenting the model’s linear predictor with an auto-covariate that is derived as a weighted function of values of the response variable in the neighbourhood of that point (Cliff and Ord 1973, Keitt et al. 2002). These spatial models exist for both normally (Keitt et al. 2002) and non-normally distributed response variables, such as the auto-logistic (Augustin et al. 1996), the auto-Poisson (Huffer and Wu 1998) or GLMMs with auto-correlated random effects (Diggle et al. 1998, Stephenson et al. 2006).

Unfortunately, likelihood estimation for the latter models requires computationally expensive high-dimensional numerical integration, which is especially problematic for the large quantity of data produced by telemetry studies.

Although the use of these spatial models will become increasingly important in future studies, the current lack of guidelines for model structure and estimation software makes them an impractical proposition. A more practical method may be to use a conservative model selection protocol involving a higher penalty in the information criteria, or to implement model selection by re-sampling, e.g. cross-validation. It should be noted that cross-validation and the use of models that explicitly model residual spatial autocorrelation are not mutually exclusive.

Problem 6: telemetry locations from the same individual are serially correlated

The degree of dependence between two successive observations (response = 1) from the same animal depends on how fast it moves, the frequency of observation and repetition in behaviour. If, additionally, the environmental variables are strongly spatially autocorrelated, the two observations are likely to occur in similar environmental conditions (Fig. 3).

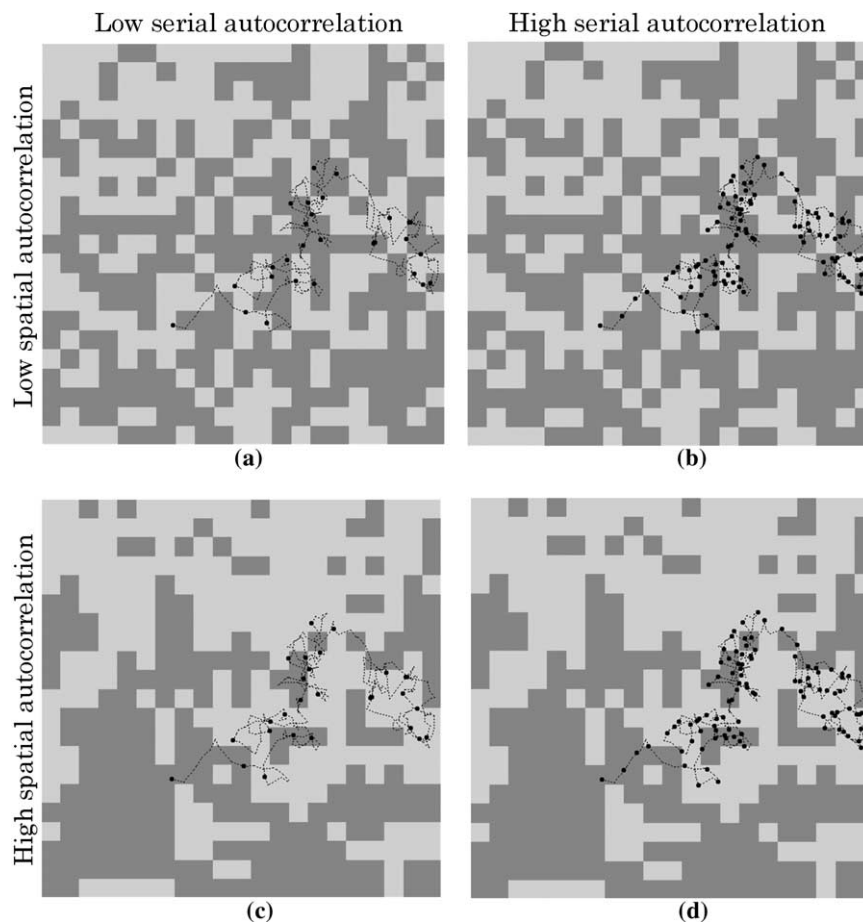


Fig. 3. The extent of inter-dependence in telemetry data is determined by how fast the animals move and how often they are observed (serial autocorrelation), and by the degree of similarity in conditions between neighbouring points in space (spatial autocorrelation). The assumption of independence is most severely violated when frequently sampled data are regressed against strongly autocorrelated environmental variables (part d).

Biologically, this may exaggerate the impression of preference for these conditions. Statistically, it will lead to underestimation of parameter uncertainty (a false increase in precision) and the possibility of retention of irrelevant environmental variables during model selection (loss in parsimony).

One solution is to discard observations based on the so-called “time to independence” (Swihart and Slade 1985). However, this treats dependence between two points as either present or absent. In reality, there are degrees of independence and censoring the data leads to a data-set that is not completely free of autocorrelation and poorer in information since the discarded points contained some information (Rooney et al. 1998). Alternatives are post-hoc adjustments of the model’s degrees of freedom or permutation tests (Legendre 1993).

Aebischer et al. (1993), point out that the objective of most ecological studies is to draw inferences about the population, and therefore that biological hypotheses must be tested at the level of the individual animal rather than the telemetry observation. Issues with temporal autocorrelation are therefore thought to be avoided by using the animal as the sampling unit.

Mixed-effect models, can explicitly capture both within-individual stochasticity (as the variance of the fixed effects) and between-individual variation (as the random effects). Some of our simulations indicate that, in mixed-effects models, residual serial correlation still causes under-estimation of the fixed effects variance, but leads to increases in the variance of the random effects. This is because the serial correlation results in artificially low, within-individual variability, which makes individuals appear more distinct from each-other. Using hypothesis-testing to make population level inferences based on the random effects, rather than standard errors of fixed-effects models, is therefore more conservative.

To draw inferences about the distribution of a species, model selection is increasingly being used in preference to hypothesis-testing. The ICs used may lead to unnecessarily complex models, if serial autocorrelation is not modelled in the likelihood. However explicitly modelling the serial autocorrelation in the data is computationally very demanding.

When modelling habitat preference, serial dependence is aggravated by the existence of strong spatial autocorrelation in the covariates (Fig. 3), the use of autoregressive or spatial error models for treating the effects of spatial autocorrelation (see above) could also help reduce the consequences of serial autocorrelation (Augustin et al. 1996). This would involve specifying the correlation as a function of spatio-temporal displacement between telemetry observations but, to our knowledge, no studies have implemented these ideas for GLMMs.

In the absence of the technical and computational capability to model serial autocorrelation for large and complicated datasets, it is imperative to find a practical treatment of its consequences for model selection. We advocate the use of cross-validation, on the basis of arguments similar to those presented for spatial autocorrelation (see Problem 5 above).

Problem 7: animals are not equally detectable in different habitats

If telemetry observations are received at a constant rate, their number in each habitat is an unbiased estimate of the proportion of time spent there. However, the rate of data acquisition may be affected by 1) behaviour (e.g. different detectability of individuals when travelling or foraging), 2) the environment (e.g. reduced signal transmission caused by dense forest canopy), and 3) satellite reception (e.g. orbital variability in satellite coverage). Environmental conditions can affect detectability and bias the estimates of usage either directly (e.g. forest canopy) or indirectly (i.e. change of behaviour).

Frair et al. (2004) suggest quantifying the rate of data acquisition as a function of environmental variables and incorporating this into the habitat preference model. However, it is difficult to obtain estimates of detection probability independently of the telemetry data, particularly for inaccessible (e.g. marine) areas and, generally, impossible to account for behavioural influences on detection.

Alternatively, the path of the individual can be reconstructed using either interpolation or smoothing techniques, to obtain a sample of locations at regular time intervals. Although this reduces the bias in parameter estimates, the precision with which the position of the animal can be obtained from a reconstructed path, at any given instant, varies with the number of observations around that instant. Consequently, the reconstructed response data in poor-reception habitats or cryptic modes of movement will be less precise (Patterson et al. in press). On the other hand, path reconstructions based on smoothing can improve overall precision by correcting some of the erroneous outliers in the raw data.

Problem 8: animals respond non-linearly to their environment

Ecologists are interested in the direction (preference/avoidance) of the animals’ response to environments but, also, in its shape. Although GLMs can include non-linear transformations of covariates, it is often unknown a-priori what these functional relationships should be. Additionally, higher order polynomials may lead to extreme (high or low) values of the transformed explanatory variable, which complicates parameter estimation and may lead to oscillatory behaviour. In generalized additive models (GAMs), the appropriate functional form can be dictated by the data using scatter-plot smoothers or splines (Hastie and Tibshirani 1990, Wood 2006). Compared to non-linear transformations, splines are on average more flexible (per DF) and they permit local flexibility through knot selection. The GAM equivalent of the linear predictor in eq. (7) is:

$$\eta_i = b_{0,m} + s_1(x_{i,1}) + \dots + s_j(x_{i,j}) + \dots$$

where

$$s_j(x_{i,j}) = \sum_{k=1}^K b_{k,j} f_{k,j}(x_{i,j}) \quad (10)$$

Smoothers are classified as either parametric (e.g. b-splines, natural splines) or non-parametric (e.g. running mean, bin and kernel). Most parametric smoothers apply a set of

pre-defined (e.g. cubic polynomial) transformations $f(\cdot)$, known as basis functions or the column of a spline, to an explanatory variable. Each basis function is constructed from the explanatory variable and a set of pre-specified points on the x-axis, known as knots, which are often based on quantiles of the explanatory variable (Ramsay 2005). Each basis function is specified by a different set of knots and therefore behaves differently at different parts of the range of values taken by the explanatory variable. The set of basis functions, evaluated at the covariate values, can be implemented as a new set of covariates, replacing each of the original environmental variables in eq. (7).

In a mixed-model approach, the b 's are random variables rather than parameters (eq. (7) and (9)). Since each basis function, applies to a particular range of the covariate, the mixed-effects approach models the amount of individual variation at different values of the covariate. This detects whether different individuals are affected by particular covariates but, also, whether the shape of this relationship differs between individuals.

Case-study: grey seals *Halichoerus grypus* in the UK

Grey seal natural history

An estimated 130 000 grey seals inhabit the coasts and seas around the British Isles (SMRU 2005). They stay on land

for breeding (Oct.–Nov.) and moulting (Jan.–March). In the rest of the year, individuals aggregate on coastal sites, known as haulouts. When at sea, grey seals perform predominantly (88% of times) return trips to the departure haulout site each lasting, on average, 2.33 d. They are not completely site-faithful, occasionally performing trips to haulouts hundreds of kilometres away (McConnell et al. 1999).

Grey seal diet varies considerably between individuals, across space and seasons, partly due to spatio-temporal variation in the availability of their >20 prey species (Hammond et al. 1994). Nevertheless, they consistently take sandeels *Ammodytes marinus*, a small cryptic species that spends part of its life buried in coarse sediment (Wright et al. 2000).

Methods

Response variable

In the period 1991 to 2001 a total of 58 grey seals were caught at the Farne Isles (55°38'N, 1°37'W), Abertay (56°24'N 3°05'W) and Isle of May (56°19'N, 2°56'W) haulout sites (Fig. 4a) and fitted with a satellite relay data loggers (SRDLs) (McConnell et al. 1999).

During their lifetime, the SRDLs sent UHF signals to two polar-orbiting Argos satellites. The location of the animal was determined using the Doppler shift in frequency of the signal (Argos 1989). These estimates




Fig. 4. Examples of data used in the model: filtered telemetry observations superimposed on the accessibility surface (a), control locations drawn from the accessibility surface (b), plot of % mud in the sediment (c) and sea-bottom depth (d). The white circles in (a and d), represent Abertay, Isle of May and the Farnes (from N to S).

are vulnerable to bias and imprecision when they are based on a low number of successive uplinks. In particular, because the distribution of the Argos observation error for poor-quality locations is thick-tailed (Vincent et al. 2002) the data contained a small number of highly erroneous location fixes which had to be removed (as per McConnell et al. 1992). To treat error in the remaining data we used a smoothing algorithm developed (M. Lonergan pers. comm.) within the MGCV package in R (Wood 2006). This uses a GAM to produce a smooth path in space as a parametric function of time. Location quality (LQ), is used to weight the influence of different locations by the inverse of their error variance, as given in Vincent et al. (2002). From these smoothed paths we obtained interpolated positions at 3-h intervals.

Habitat preference depends on the type of activity performed by the animals. Different habitats may be suitable for different activities such as resting, breeding, migrating or foraging. We were primarily interested in foraging behaviour and therefore restricted our attention to telemetry data from return trips (individual forays outside the 5 km radius around the haulout). We only included individuals for which we had in-excess of 7 d of return-trip data. We used the Farnes data for fitting/model-selection and the Abertay data to assess the model's ability to extrapolate geographically.

The error-corrected and temporally-interpolated animal locations represented a sample of points in geographical space visited by the tagged grey seals (presence data). These were complemented with a set of points not visited by the study animals (control data, i.e. response = 1). Some control points could, in fact, be used by the animal. Since each matches to a telemetry location, we can assume that the animal did not visit that point in space at that time instance. Control points were selected from the surface of accessibility calculated by Matthiopoulos et al. (2004)

$$\alpha = \left(\frac{d}{5}\right)^{-1.98}, \quad (11)$$

where d is the at-sea distance (in km) between any point in space and the haulout.

The mathematical results pertaining to case-control studies (Prentice and Pyke 1979) are asymptotic, meaning, in this context, that habitat availability is only approximately represented in the sample of response data, the approximation improving with an increasing number of absence data. The number of absence data necessary to obtain a sufficiently good approximation will vary from study to study. For our data, estimated coefficients remained effectively unchanged when the 0s were at least twice the 1s.

Explanatory variables

We selected environmental variables on the basis of possible biological relevance and the availability of data. In studies of habitat preference, the variables that are seen as most relevant to the study-species are often the most data-poor. In apex marine predators, such as grey seals, this predicament is particularly true because grey seals eat a large number of prey species, each having its own patterns

of spatio-temporal variation and because fish distribution data are notoriously difficult to collect. We found little readily usable prey data and opted instead for three static variables: sediment type, sea depth and distance from the haulout. Although this decision was motivated by data sparsity, it was further justified by the requirements for a parsimonious, predictive and, yet, biologically relevant model: a model using static environmental variables is more parsimonious because it only requires that the distribution of foragers is the result of time-invariant environmental cues. Furthermore, developing a predictive model of grey seal distribution that relied on the continued measurement of species that are just as variable would defeat the purpose of the exercise. So, even if prey data were available, they would be of little use for future model predictions. On the other hand, if prey distributions are predictable, either by seals or by human observers, it will partly be from cues provided by geophysical variables such as the ones we have used.

Sediment type was derived from British Geological Survey (BGS) data obtained from core samples, spaced, on average, at 5 km intervals throughout the study area. The raw data were given as a percentage-by-weight of gravel (defined as particles >2.0 mm in diameter), sand (particles 0.0625–2.0 mm in diameter) and mud (particles <0.0625 mm in diameter). We used a random sub-sample of cores to calculate the semi-variogram (Isaaks and Srivastava 1990) characterizing the spatial autocorrelation of each of the three sediment components. The semi-variograms were then used to generate kriged estimates of each sediment component throughout the study region, at a resolution of 1 km². We kriged the three sediment components independently despite the fact that they were complementary percentages but the sum of estimated percentages over all grid cells had an average value of 101% with standard deviation of 5%. Local estimates were then normalized to 100%. For modelling, we used only two (% mud and gravel – Fig. 4c) of the three sediment components to avoid severe problems of multi-colinearity (see above).

Sea depth data was obtained from the British Geological Survey (BGS) Digbath250 <<http://www.bgs.ac.uk/products/digbath250/>>. This was represented by bathymetric contour lines every 10 m between 0 and 100 m, every 20 m between 200 and 400 m and every 100 m at depths exceeding 400 m. We transformed these maps using an equal-distance projection (National Grid of Great-Britain). We placed points at 1 km intervals along the contour lines and interpolated linearly between these points to obtain depth estimates for every point of the regular, 1 km grid (Fig. 4d). To account for residual issues of accessibility, we used at-sea distance from the seals' departure haulout calculated on a grid of 1 km resolution (Fig. 4a).

Model structure

The basic structure of the model followed from eq. (3)

$$\hat{u}_i \sim B(1, h_i) \\ h_i = g^{-1}(\eta_i) = \frac{e^{\eta_i}}{1 + e^{\eta_i}}. \quad (12)$$

The linear predictor (η_i) was structured as a GAM with a maximum of four covariates

$$\eta_i = b_0 + s_1(\text{distance}_i) + s_2(\text{depth}_i) + s_3(\text{mud}_i) + s_4(\text{gravel}_i) = b_0 + \sum_{k=1}^6 b_{1,k} f_{1,k}(\text{distance}_i) + \dots, \quad (13)$$

where each spline $s(\cdot)$ was a flexible function of each covariate, and $b_{j,k}$ was the random effect parameter for the k th out of 6 basis function of the spline for the j th explanatory variable. We used b-splines because they afford local control in modelling the response at different domains of the explanatory variable and produce robust results in data-poor regions of environmental space (Ramsay 2005). For each explanatory variable, we used a b-spline (with intercept) requiring 6 DF; two boundary knots and two internal knots at the 33 and 66% quantiles of the frequency distribution of observed values of the environmental variable at the coordinates of the telemetry observations. Although it is possible to select spline flexibility by cross-validation (Wood 2006), computational restrictions forced us to pre-specify the position and number of knots. However, as a minimum preventative measure against over-fitting, we compared the performance of the GAM with the output of a GLM.

Equation (13) was implemented as a mixed-model. Each parameter $b_{j,k}$ was treated as a normal variable, containing a fixed effect $\beta_{j,k}$ (applicable to the entire population of animals), and a random effect error term $v_{j,k} \sim N(0, \sigma_{j,k})$, representing the variability in the response of different animals to the k th spline of the j th explanatory variable. This model is similar to a generalized additive mixed model (GAMM – Wood 2006). Individual variation may partly be due to individual characteristics such as sex and length,

$$b_{j,k,m} = \beta_{j,k,0} + \beta_{j,k,1} \text{sex}_m + \beta_{j,k,2} \text{length}_m + w_{j,k}, \quad (14)$$

where m refers to an individual animal and “length” was measured from nose to tail. Incorporating these helped account for possible biases in sampling effort across different types of animals and was a practical way of modelling deviations from normality in the random effect (see above). Other possibilities were the mass and age of the animals. We excluded mass because it was strongly correlated with length and it is harder than length to observe remotely. Age was excluded because it wasn’t recorded for 5 individuals and its inclusion would necessitate a considerable reduction in sample size.

Parameter estimation

GAMMs are usually estimated by penalized quasi-likelihood (PQL) or maximum-likelihood methods (ML). PQL is fast and produces nearly unbiased parameter estimates but does not produce a likelihood for use in model selection. ML techniques produce exact marginal likelihood estimates, but are computationally expensive because, for every candidate set of parameter values, they need to perform a numerical approximation of the likelihood over all possible realizations of the model’s random effects term. The likelihood function for mixed-models is (Pinheiro and Bates 2000)

$$L(\beta, \psi | \hat{u}) = \prod_{m=1}^M \int p(\hat{u}_m | b_m, \beta, \psi) p(b_m | \psi) db_m, \quad (15)$$

where b_m is a vector of random effects, each applying to a k th basis function for the spline of the j th covariate (eq. (14)) of

the m th of a total of M individuals. The integrand in eq. (15) consists of the response and the random effect components. The random effects component ($p(b_m | \psi)$), is a multivariate normal probability density function with means 0 and variance-covariance matrix ψ . The random effect variances (diagonal elements of ψ) represent the individual variation in the animals’ response to environmental variables. The covariances (off-diagonal elements) quantify within-individual correlations in the response. For GAMMs, the number of covariances is generally large and estimating ψ is computationally difficult. We therefore only estimated the variances of the random effects.

The response component of the likelihood was modelled as a binomial probability

$$p(\hat{u}_m | b_m, \beta, \psi) = \prod_{i=1}^n h_i^{\hat{u}_{m,i}} (1 - h_i)^{(1 - \hat{u}_{m,i})}, \quad (16)$$

where $\hat{u}_{m,i}$ is i th observation from the m th individual and h_i is the response as defined in eq. (12–14). The ML estimates are the values of the parameters β and ψ , that maximize eq. (15). We used the automatic-differentiating model-builder (ADMB) and its random effects module (Skaug 2002, Otter Research 2004) to minimize the quantity $-L(\beta, \psi | \hat{u})$. ADMB-RE first approximates the likelihood using the Laplace approximation and importance sampling. It then uses automatic differentiation to obtain exact derivatives of this approximation, which it maximizes with a quasi-Newton method (Skaug 2002).

Model selection and model validation

Information criteria (IC) such as the Akaike information criterion (AIC) or the Bayesian information criterion (BIC) (Burnham and Anderson 2002) achieve parsimony by penalizing the likelihood of the model by the number of parameters it contains. Therefore, the effectiveness of IC relies on the correctness of the assumptions underpinning the likelihood function. In telemetry studies, violation of the independence assumption can lead to over-fitted models (see above). We used IC only as a rough guide and relied on cluster-level cross-validation for final model selection. We initially fitted GAMMs to a subset of the data (19 out of 29 individuals) from the Fame Isles. When defining the model and validation datasets, we split the different types of individuals equally to facilitate the detection of individual-specific effects on habitat preference. We started with an intercept-only model and used forward model-selection (with AIC) to arrive at a model containing all four environmental variables. This reduced the number of models to be investigated by cross-validation from a possible 41 to 5. From these, we selected the one with the highest explained deviance in predicting the data from the remaining 10 animals. This still left room for over-fitting because the number of knots used for each environmental variable was pre-specified. As a minimum measure against over-fitting, we replaced each spline in turn by a linear term and retained the changes that improved the model’s performance in cross-validation.

The selected model was then extended with individual characteristics (i.e. sex and body-length). Again, forward model selection based on AIC was used to arrive at a set of candidate models, each of which was validated against

the independent test data set to yield a final model. To assess this model's predictive performance, we estimated the likelihood of our model and its parameter estimates, under data from 13 individuals from a neighbouring sub-population, in Abertay. We compared this value with, similarly calculated, likelihoods for 5 other candidate models.

Predicting spatial usage

Equation (6) can be used to calculate the expected usage $f_u(X_s)$ around every point s in space, given values for the environmental variables and accessibility $f_a(X_s)$. This requires predicting h_i which, for random-effects models, is not straightforward: assuming, for simplicity, that individual characteristics are not included in the model, the random effects b_j are normally distributed with mean $\bar{b}_j = \beta_j$ and variance σ_j^2 . The response h_i is modelled as a non-linear function (logit^{-1}) of the linear predictor containing the random effects. This means that the distribution of the random effects viewed on the scale of the response is transformed and, therefore, the fixed effects part β_j is no longer the mean of this distribution. Consequently, predictions of usage generated from the model using the fixed effects alone, are not the same as the average predictions from multiple realizations of the random effects. In practice, the best way to calculate h_i , is as the average of a random sample of predictions generated from the estimates for the fixed effects β_j , incremented by a value drawn from $N(0, \sigma_j^2)$.

Individual characteristics (e.g. sex and body length), can be drawn randomly from the pool of observed values associated with the sample of tagged individuals but, to avoid catching effort biases, it is better to use other sources. For example, the age-/sex-structure could be obtained from demographic models (but this requires precise estimates of vital rates or age structure) and the distribution of body lengths could be derived from aerial transect surveys.

We predicted usage $f_u(X_s)$ on a 1 km resolution grid. Every such map obtained for a single realization of the model's random effects can be thought of as the space-use of one individual from the population. However, such surfaces did not exactly add up to 1 (and needed to be normalized) primarily for two reasons: first, the environmental conditions at some points in space were outside the range of the data used to fit the model (e.g. the distribution of values for the sediment and depth variables was heavily skewed to the right). Extrapolating usage to these points occasionally led to spurious, local over-estimates. We therefore excluded predictions outside the ranges of these environmental variables. Second, only a sample of points was included in the analysis and therefore, by chance, relatively more- or less-preferred environmental conditions might exist within the study area compared to the data.

Some applications may require the absolute, rather than relative, population density in space. Matthiopoulos et al. (2004) estimated that there are 2950 (95%CI: 2742–3159) individuals associated with the Farne Isles. Additionally, our analysis disregarded 1) usage within 5 km of the haulout, 2) regions in space for which no environmental data were available and 3) all points that were outside the range of environmental data. Expected population density was

therefore approximated by scaling usage by $N(1-p)$ where p is the proportion of telemetry observations that fell in the excluded spatial regions.

Predicting preference

The areas receiving more usage than expected by chance (taking into account accessibility), are said to be preferred. Spatial plots of preference indicate the position and extent of these areas, also known as “hotspots”. Preference w at a point s can be calculated using eq. (7) which holds for all r . As the number r of controls per observation increases, the response variable h_i becomes small, usually much smaller than 0.5. For these values of h the ratio $h_i/(1-h_i)$ is approximated well by h ,

$$w = r \frac{h_i}{1-h_i} \cong h_i r \propto h_i \quad (\text{for large } r), \quad (17)$$

which is in agreement with eq. (S4) (Supplementary material) derived for r tending to infinity. Equation (7) relies on accurately modelling accessibility. To account for biases in $f_a(X_s)$ we included distance to the haulout as an additional covariate. Hence, plotting $w(X_s)$ from eq. (7), does not purely reflect the animals' preference for environmental variables. By breaking up the linear predictor into three additive components, preference can be re-written

$$\frac{f_u(X_s)}{f_a(X_s)} = \text{rexp}(\beta_0) \exp(\eta_{\text{distance}}) \exp(\eta_{\text{environmental variables}}), \quad (18)$$

which implies

$$\frac{f_u(X_s)}{f_a(X_s) \exp(\eta_{\text{distance}})} \propto \exp(\eta_{\text{environmental variables}}). \quad (19)$$

The expression on the left is preference, corrected for the biases in the accessibility model, and therefore $\exp(\eta_{\text{environmental variables}})$ can be used to generate spatial plots of preference.

Results

Response and environmental variables

From a total of 58 individuals caught on the Farnes, Abertay and the Isle of May, 42 provided enough data to be included in the analysis. Tag life (minimum 2.5 d, maximum 329 d, average: 109 d) and transmission intensity (minimum 0.91, maximum 10.4, average: 5.9 locations d^{-1}) varied greatly between individuals, leading to large differences in sampling intensity between individuals. Most individuals were caught in early spring and summer and consequently there was large temporal overlap in sampling intensity between individuals, with most data around August and the least in February. Preference may change due to temporal changes in resource availability, the environment, or individual requirements. This implies that our results may only apply to the period of data collection. Following pre-processing of the satellite data (error-correction and path reconstruction) the Farnes data-set comprised a total of 2315 telemetry observations (Fig. 4a). We randomly selected twice as many points from the accessibility surface (Fig. 4b). The combined

presence and absence data constituted the response data that were regressed against the environmental variables (examples shown in Fig. 4).

Model selection

Model-selection (Fig. 5a) using the Farnes data suggested a model containing all environmental variables. The variables

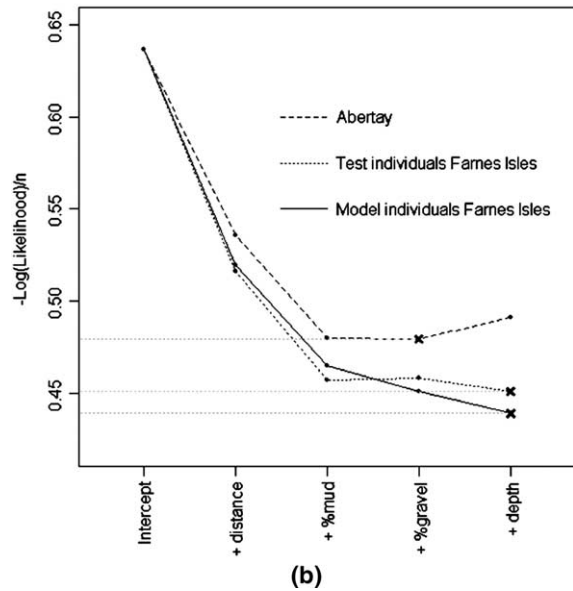
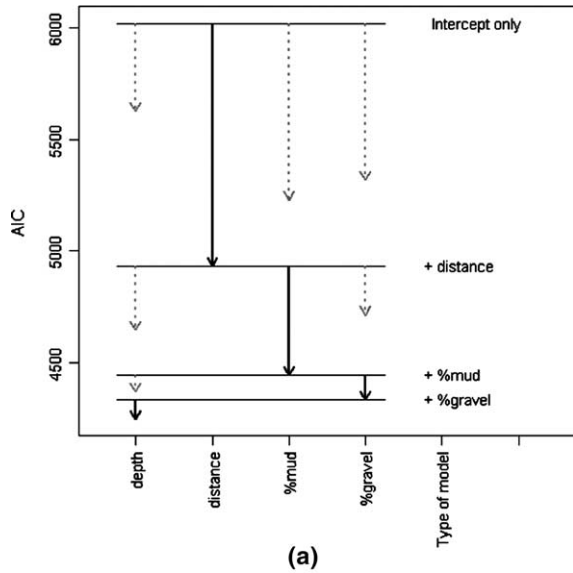


Fig. 5. Model selection (a) and validation (b) using only environmental covariates. Forward model selection based on 19 individuals from the Farnes (a) resulted in a sequence of increasingly complex models. The length and direction of the arrows indicate the change in AIC following the addition of each covariate. Solid arrows indicate the variables that led to the largest improvement in AIC. Using cross validation with additional data from 10 different animals of the same sub-population (b), enabled us to select the most parsimonious model from the sequence of models generated in (a). Using data from the Abertay sub-population enabled us to validate the robustness of model selection under extrapolation. A black cross indicates the best model for each of the three data sets.

“distance to the haulout site” and “% mud in the sediment” accounted for 87% of the total explained deviance. This model also performed best with the Farnes validation data set (Fig. 5b). The sequential replacement of the splines by un-transformed environmental variables generally led to inferior models. However, a model with no splines for gravel performed better with both validation data sets.

Based on the AIC, the interactions between %mud and body length, %mud and sex, distance to the haulout and

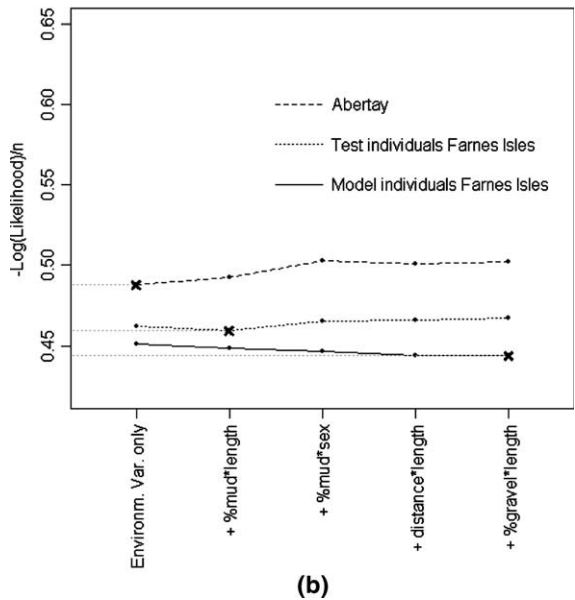
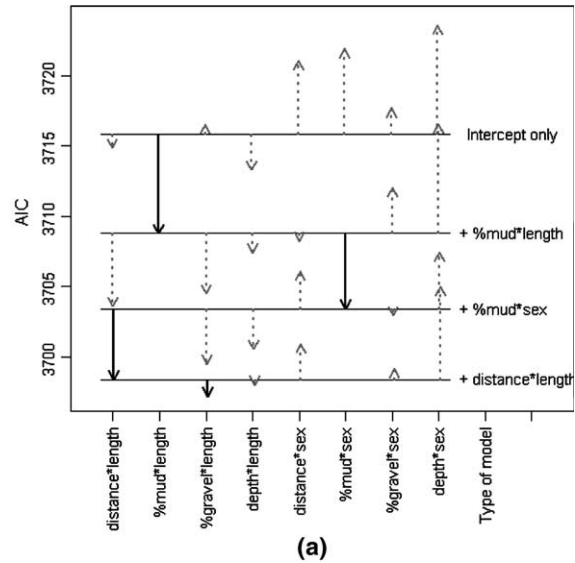


Fig. 6. Model selection (a) and validation (b) for models containing individual characteristics. Starting with the model selected in Fig. 5, we incrementally added individual characteristics and interactions (a). Additional data from different animals belonging to the same sub-population (Farnes Isles), enabled us to select, from the candidate models created in (a), a parsimonious model by cross-validation (b). Data from the Abertay sub-population were used to validate the robustness of model selection under extrapolation. A black cross indicates the best model for each of the three data sets.

length, and %gravel and length led to better models (Fig. 6a) but, only the interaction between the nose-to-tail length of the individual and %mud in the sediment led to an improvement (Fig. 6b) under cross-validation.

Assessment of predictive performance using the Abertay data indicated that the model containing all environmental variables was outperformed slightly by a model without depth. Inclusions of individual characteristics did not improve the predictive performance.

We used semi-variograms to detect spatial autocorrelation in the response and residuals. The response variable was positively autocorrelated but the residuals were not, suggesting that all important covariates were included in the analysis.

Relationships of environmental variables with the response

Distance to the haulout had a tight negative relationship with the response, implying that usage far away from the haulout was over-predicted by the accessibility model (Fig. 7a). The second most important variable was %mud which also had a significant interaction with body length. The animals' preference was highest for areas with mud content of ca 3% (Fig. 7b). While small individuals (Fig. 7c) changed their response relatively little with increasing mud-contents, large animals (Fig. 7d) had a well-defined peak in preference for those areas. Our grey seals also preferred gravelly areas (Fig. 7e) and appeared to have an increasing preference for depths down to 80 m,




Fig. 7. Effect of the covariates distance to the haulout (“distance”) (a), the interaction between “%mud” and “nose-tail length (cm)” (b), “%mud” for small (c) and large (d) individuals (“nose-to-tail length” is 100 and 180 cm, respectively), “%gravel” (e) and “depth” (f), on the response variable on the scale of the link-function (y-axis). Shaded areas represent 95% confidence limits, the gray lines represent individual-specific responses and the black line gives the mean population responses (i.e. the fixed effect).

after which mean preference decreased slightly. This tendency was confounded by large individual variability (Fig. 7f).

Spatial prediction of usage and preference

We plotted estimated usage $\hat{f}_u(X_s)$ for every point in space with available estimates of environmental conditions (Fig. 8a). Since it is often difficult to visually compare the raw telemetry data with a density plot we also plotted a realization of a heterogeneous spatial Poisson process with

rate $f_u(X_s)$ (Fig. 8b). We used the model fitted to the Farnes data, to predict usage for Abertay seals (Fig. 8c, d). To indicate which areas grey seals would use if all points in space were equally accessible, we plotted estimated preference in space (Fig. 8e).

Discussion specific to grey seals

All the environmental variables examined in this case study accounted for some variation in the response. Distance to

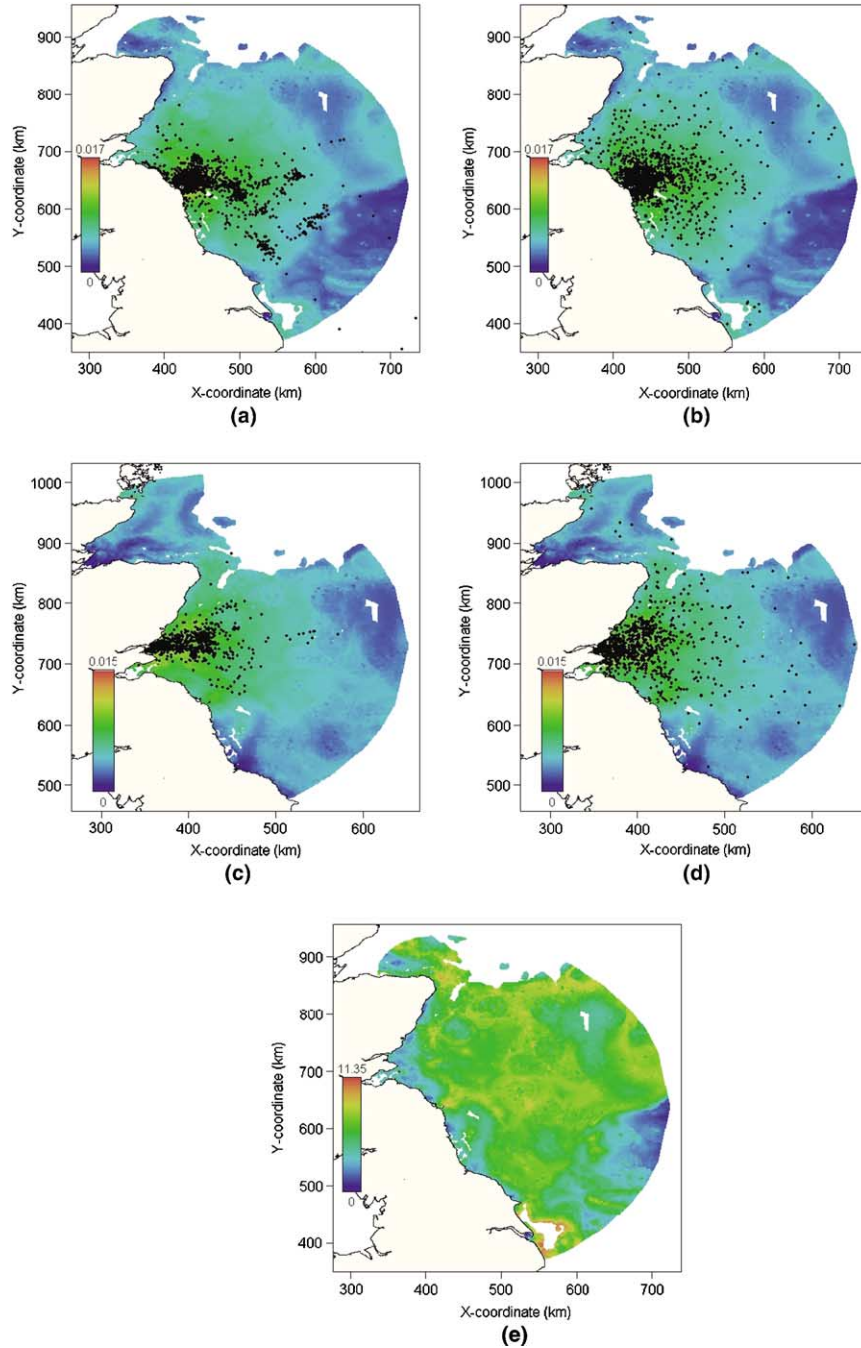


Fig. 8. Plots of goodness-of-fit (a and b) and extrapolation (c and d) of usage and preference (e): spatial prediction of usage for the Farnes Isles showing the telemetry observations (a) and one possible realization from the model using the same sample size (b). Spatial prediction of usage for Abertay using the Farnes Isles model showing data (c) and a realization from the model (d). Spatial prediction of preference corrected for unequal accessibility (e).

the haulout and %mud were the most important. The negative relationship between the response variable h and distance indicates that seals concentrate closer to the haulout than is predicted by the accessibility model. Avoidance of mud may result directly from the seals' bio-turbating search tactics (Bowen et al. 2002) or indirectly from the distribution of sediment-burrowing prey such as sandeels (Wright et al. 2000).

Depth might have been expected to be an important covariate because seals are predominantly benthic feeders whose useful foraging time at the sea bed decreases with increasing depths. However, the entire North Sea is relatively shallow and therefore unlikely to test the grey seals' diving capability. A second reason why depth might have been expected to be important is its strong correlation with the density of sandeels. Wright et al. (2000) have shown that most sandeels around the Shetland Isles, in the North of Scotland, occur around depths of 50–60 m. However, this relationship does not necessarily hold across the North Sea. Also, if sandeels occurred in relatively deeper areas around the Farnes Isles, this might counter the additional cost of diving to such depths.

General discussion

Recognition of the importance of space in population dynamics (Tilman and Kareiva 1997) and ecological interactions (McLaughlin and Roughgarden 1992) has greatly advanced the development of data-collection techniques on species distribution. For transect data, this increase in activity has been matched by the development of analytical techniques (Buckland et al. 1993). Less effort has been directed at telemetry data primarily because none of the off-the-shelf statistical frameworks can address all the problems they pose but, also, due to the historical confusion in the interpretation of model components as crucial as the response variable.

Hence, many applied studies proceed with the analysis of telemetry data without acknowledging these problems and ambiguities. Although this may not always affect their results, there is nevertheless a distinct risk that conservation and management of some populations is being based on false premises. In this paper, we 1) highlighted these conceptual ambiguities and proposed a self-consistent framework of definitions for terms old and new, 2) provided a summary of the evolution of regression analysis of space-use and habitat preference, 3) enumerated, described and proposed solutions for the eight fundamental challenges of analyzing telemetry data and 4) provided an illustration using a case study.

Our framework brings together case-control, mixed-effects and generalized-additive modelling. We believe that this combination best addresses most of the problems listed in the first half of the paper and is also flexible enough to benefit from forthcoming advances in computer speed, estimation software and statistical methodology to address those problems (e.g. spatial and temporal autocorrelation) that were not tackled definitively here.

Empirical models of usage and preference

The presentation of our framework was problem-oriented. However, it is now useful to collect our insights into the fundamental components of empirical models: response and explanatory variables, model structure, parameter estimation, model selection, validation and prediction.

Response variable

Obtaining a response variable from wildlife telemetry data is challenging because they only provide information about animal presence. Some studies (Drake et al. 2006) have regressed the density of telemetry observations against environmental covariates, a technique known as niche-based analysis. This makes the stringent assumption that sampling effort is uniformly distributed in space and ignores variations in the relative availability of different habitats. In practice, even if an animal showed no preference, it would still be observed more frequently in more abundant habitats. This has traditionally been resolved by classifying environmental space into habitats and correcting usage for their relative availability (Manly et al. 2002).

Our work and other, related, studies (Boyce and McDonald 1999), have replaced classification in geographical and environmental space by a case-control design, allowing the model to capture the finest variations in usage, within the limits of computation. This causes ambiguity in the interpretation of the response variable in the resulting models. Keating and Cherry (2004) have shown that the response variable in a case-control design is not proportional to usage. Here, we have shown how usage can be estimated from it.

Confusion also exists about the positioning of controls for a case-control design. The controls can be selected uniformly randomly from the postulated range of the animals or from a surface of accessibility as was done here. Given a sufficient number of controls and including distance to the central place as a candidate covariate, different sampling designs $f_a(X_s)$, will produce comparable estimates of usage $f_u(X_s)$ (eq. (6)), but a model of accessibility is preferable for large data sets because it directs more computational power to those areas that are more likely to be visited by the animals.

Explanatory variables

Using all conceivable candidate covariates in regression can be “a subjective and iterative search for data patterns and significance”, and model selection should, instead, be conducted from a small set of covariates, believed to be biologically related to the response variable (Burnham and Anderson 2002). However, we argue that excluding particular combinations of covariates a priori, is no less subjective. E.g. in our grey seal example it might have been argued that sediment variables should only be considered as proxies for prey distribution since “seals don't eat sediment”. However, experiments with captive animals (SMRU unpubl.), have indicated that seals can forage by continuously disturbing the sediment with their snout causing fish to appear. Seals may therefore actively

choose to forage in sandy regions to enhance their consumption rates. So if, contrary to biological intuition, a covariate is retained by out-competing others in the model selection process, this may hint at an ecological process not previously considered and thus enhance our understanding of the mechanisms underlying the species' distribution. Given the exploratory nature of most current telemetry studies, it is perhaps better to examine as many candidate covariates as permitted by sample size and computer power, and allow issues of parsimony to be dealt with by model selection.

Model structure

The binary, case-control data were modelled using a logit link with a Bernoulli likelihood. The hierarchical structure of the data, was represented by a mixed-effects model (Pinheiro and Bates 2000) which provided an estimate of individual variation, necessary for population-level inferences. In contrast to predictions based on transect data, this can reveal whether an animal observed at extreme environmental conditions (e.g. a grey seal off the continental shelf) is a one-off event for that individual or a one-off individual for that population. Furthermore, although capture of animals is often biased towards certain types of individuals (e.g. young males), mixed-effects models can quantify the contribution of individual characteristics on the observed variation. Independent information on population structure can then be used to generate balanced predictions of population distribution.

The advantage of GAMMs is that they can highlight differences in the functional form of the responses of different individuals (e.g. Fig. 7). For example, our GAMM indicated that grey seal preference peaks at depths of ca 80 m. Due to the dominance of data close to shore, at shallower depths, a GLM would have unrealistically predicted a continued positive trend. So GAMMs generally increase the model's predictive capability, if non-linear trends are not caused by overfitting, which can be prevented with careful model selection.

Model selection and model validation

Instead of testing for the significance of candidate covariates at some arbitrary threshold α , model selection offers an objective way of ranking their explanatory power (e.g. Fig. 5a, 6a). Model selection is often implemented using information criteria (IC) which penalize the likelihood of a candidate model by the number of parameters it contains. Although there are theoretical justifications for the severity of the penalty (Burnham and Anderson 2002), their validity is sensitive to miss-specification of the likelihood function. For computational reasons, the likelihood of most non-linear models assumes that the data are independent. When, as with telemetry studies, the data are spatially and temporally autocorrelated, model selection by IC leads to over-fitted models (Burnham and Anderson 2002). This can be overcome either by modelling autocorrelation as part of the likelihood, or by using alternative approaches, such as cross-validation, for model selection. Cross-validation prevents over-fitting by using one data set for fitting the model and another for assessing its predictive power. In our case-study, cross-validation led to a reduction of the number of

environmental and individual-specific covariates, and improved spatial predictions, compared to the models suggested by the IC.

Predictions of usage and preference

Spatial predictions can be classified into three categories in order of decreasing reliability; 1) predictions for the area and time in which the telemetry data were collected (interpolation), 2) predictions from a different place or time, but for similar environmental conditions as those used to construct the model (geographical extrapolation) and 3) predictions outside the region of environmental space for which data exist (environmental extrapolation).

Interpolation is subject to biases in catching effort, large individual variation and a small sample size. Mixed-effects models with individual-specific covariates can account for some of these biases and the use of cross-validation ensures that a model fitted to one group of individuals is applicable to other individuals from that same sub-population.

Differences in absolute habitat availability between different geographical regions can weaken predictions of usage (Myrsetrud and Ims 1998). Our ability to predict the distribution of Abertay animals using data from the Farnes was probably the result of similarity in the conditions experienced by these two sub-populations, implying that our predictions were extrapolations in geographical, but not environmental space.

Extrapolation in environmental space is generally less reliable and is best avoided. For example, the grey seal data collected on the east coast of Scotland provided no hint as to the maximum depth that grey seals can dive to. Applied to the west coast of Scotland, this model might predict seal usage beyond the continental shelf at depths that are known to exceed the physiological capabilities of the species.

Wider context

Empirical vs mechanistic modelling

We employed flexible regression models (GAMMs) because they are ideal for describing unknown, non-linear relationships. We also used information from seal natural history to provide post-hoc interpretations of the predicted effect of particular covariates on the species' spatial distribution. We chose this, primarily empirical, approach because our priority was on statistical inference and because, as with many other telemetry studies, there was insufficient scientific knowledge to enable us to construct a more mechanistic model.

Generally, this decision carries two penalties: first, the outcomes of model-fitting and selection can be sensitive to stochasticity, hampering comparisons between published inferences from different studies and, second, environmental extrapolation is based on pattern instead of process and is therefore unreliable.

For some species, it may be possible to construct more mechanistic models describing some of the causal relationships between space-use and its covariates. For example, in grey seals the results of physiological experiments might be used to construct energetic models of the cost of travelling to particular depths or distances from the haulout (Thompson and Fedak 2001). However, it is

not always clear how to embed such models in a formal statistical framework. One possibility is to use their output to construct candidate covariates (e.g. our use of the movement model by Matthiopoulos (2003b) for accessibility).

Alternatively, mechanistic models can be fitted directly to data. Traditionally, statistical inference was considered solely with reference to purely empirical models because observation error could not be modelled along with process stochasticity and because fitting such models was computationally prohibitive. These restrictions have been relaxed with the advent of more mechanistic Bayesian state-space models (Jonsen et al. 2003, Patterson et al. in press), more powerful computers and new model-fitting algorithms. Although this is a promising area of research, currently only simple models with few covariates can be fitted, model selection is computationally expensive and population-level predictions rely on individual-based simulation.

Eulerian vs Lagrangian models

The dichotomy between Eulerian and Lagrangian models is a fundamental theme in spatial ecology (Turchin 1998). Eulerian approaches focus on variations in abundance around particular points in geographical space whereas Lagrangian approaches focus on individuals as they move across different points (Morales et al. 2004, Patterson et al. in press). There is a direct correspondence between Eulerian density and transect data just as there is congruence between Lagrangian trajectories and wildlife telemetry data. Trying to shoe-horn telemetry data into a Eulerian framework, such as the one presented here, is the cause of many of the problems addressed in this paper. This was necessary because a Eulerian approach is able to process both presence and absence data and, therefore, unlike existing Lagrangian approaches for telemetry data, provides the contrast necessary to model preference. Also, Eulerian approaches are ideal for large-scale, exploratory studies with an applied agenda because they are better suited to population-level inferences than the more behaviourally-orientated Lagrangian modelling.

Hierarchical decisions in space-use and habitat preference

Our results were predicated on the study animals' capture haulouts, but their presence in these locations was, itself, a consequence of historical processes and decisions affected by environmental variables. Early work on space-use and preference had recognized the importance of multiple levels of decision-making. According to Johnson (1980) first-order selection determines the physical or geographical range of a species. Within that, second-order selection determines the home range of an individual or social group. Third-order selection pertains to the usage made of various habitat components within the home range. Finally, if third-order selection determines a feeding site, the actual procurement of food items from those available at that site is termed fourth-order selection. Our analysis, refers to third-order selection, employing a smooth model of accessibility instead of a demarcated home-range. Had individuals been sampled randomly from the population (rarely possible, in practice), their capture locations could also be used to investigate second-order selection.

Future directions; multilevel modelling

Although we only examined the effect of individual characteristics, the mixed-effect approach can be extended to account for all the variance components outlined in Fig. 2. Variation across sub-populations could be modelled as a function of sub-population characteristics, such as its average size in individuals, its geographical position, or a measure of the overall availability of resources. This extension would be particularly useful in making predictions for haulout sites for which no wildlife telemetry data exist. Similarly, variability between trips of the same individual could be incorporated as a random effect to account for different activities performed in trips of different duration.

At the lowest level, variability between single observations could be treated as a random effect. Most modelling approaches assume that the response and explanatory data are error-free but this is not usually the case. Often, the magnitude of observation error is known but not used. For example, ARGOS provides a location quality (LQ) index for each observation which can be recast into an estimate of precision. With some additional work, this can also be obtained for spatio-temporally smoothed and regularized data. Modelling these errors as part of a mixed-effects likelihood would account for more residual noise.

All of these extensions to the mixed-effect model are conceptually simple, but computationally difficult because the time taken to approximate the model's likelihood increases exponentially with the number of random effects. Such increases in model complexity may also lead to overparameterised models. Therefore, multilevel modelling would benefit from computational advances and improved guidelines for model selection.

Conclusion

Wildlife telemetry data is potentially highly informative about the nuances of individual behaviour, but cumbersome for use in population-level inference. Ironically, predictions of population distributions are the main objective of many telemetry studies and the primary concern of end-users such as conservationists and wildlife managers. Bringing modern statistical methods to bear on this problem leads to a suitable framework for analyzing telemetry data and clarifies the interpretation of, and relationship between, the important ecological concepts of space-use and habitat-preference.

Ultimately, population distribution emerges from individual movement and statistical pattern is the result of biological process. We therefore recognize the potential for exploring its link with more individual-based approaches and extending the mechanistic properties of this framework. However, given the complexities of individual behaviour and the different ecosystems in which animals live, we believe that both of these aims can be achieved by gradually augmenting regression models with new features that address the specific challenges posed by case studies, such as the grey seal example outlined here.

Acknowledgements – All R-developers, BGS and C. Graham in particular for sediment data, S. Wood, M. Lonergan, L. Thomas and C. Donovan for their invaluable feedback, I. L. Boyd for his practical support, J. Bakker for Fig. 2, H. Skaug for help with ADMB-RE. GA was supported by grant XGLD85 to L. L. Boyd from DSTL. JM and BMcC worked under a NERC core-grant to SMRU.

References

- Aebischer, N. J. et al. 1993. Compositional analysis of habitat use from animal radio-tracking data. – *Ecology* 74: 1313–1325.
- Agresti, A. 1996. An introduction to categorical data analysis. – Wiley.
- Allredge, J. R. and Ratti, J. T. 1992. Further comparison of some statistical techniques for analysis of resource selection. – *J. Wildl. Manage.* 56: 1–9.
- Argos 1989. Guide to the Argos system. – Argos CLS, Toulouse.
- Arthur, S. M. et al. 1996. Assessing habitat selection when availability changes. – *Ecology* 77: 215–227.
- Augustin, N. H. et al. 1996. An autologistic model for the spatial distribution of wildlife. – *J. Appl. Ecol.* 33: 339–347.
- Blundell, G. M. et al. 2001. Linear home ranges: effects of smoothing, sample size, and autocorrelation on kernel estimates. – *Ecol. Monogr.* 71: 469–489.
- Bowen, W. D. et al. 2002. Prey-dependent foraging tactics and prey profitability in a marine mammal. – *Mar. Ecol. Prog. Ser.* 244: 235–245.
- Boyce, M. S. and McDonald, L. L. 1999. Relating populations to habitats using resource selection functions. – *Trends Evol. Ecol.* 14: 268–272.
- Boyce, M. S. et al. 2002. Evaluating resource selection functions. – *Ecol. Model.* 157: 281–300.
- Buckland, S. T. E. and Elston, D. A. 1993. Empirical models for the distribution of wildlife. – *J. Appl. Ecol.* 30: 478–495.
- Buckland, S. T. et al. 1993. Distance sampling: estimating abundance of biological populations. – Chapman and Hall.
- Burnham, K. P. and Anderson, D. R. 2002. Model selection and multimodel inference. A practical information-theoretic approach, 2nd ed. – Springer.
- Calenge, C. 2006. The package “adehabitat” for the R software: a tool for the analysis of space and habitat use by animals. – *Ecol. Model.* 197: 516–519.
- Calenge, C. et al. 2005. K-select analysis: a new method to analyse habitat selection in radio-tracking studies. – *Ecol. Model.* 186: 143–153.
- Cliff, A. D. and Ord, J. K. 1973. Spatial autocorrelation. – Pion, London.
- Cramer, E. M. 1985. Multicollinearity. – In: Kotz, S. et al. (eds), *Encyclopaedia of statistics*. Wiley, pp. 639–643.
- Diggle, P. J. et al. 1998. Model-based geostatistics. – *Appl. Stat.* 47: 299–326.
- Drake, J. M. et al. 2006. Modelling ecological niches with support vector machines. – *J. Appl. Ecol.* 43: 424–432.
- Fedak, M. 2004. Marine mammals as platforms for oceanographic sampling: a “win/win” situation for biology and operational oceanography. – *Mem. Nat. Inst. Polar Res.* 58: 133–147.
- Fedak, M. et al. 2002. Overcoming the constraints of long range radio telemetry from animals: getting more useful data from smaller packages. – *Integr. Comp. Biol.* 42: 3–10.
- Fox, J. 1997. Applied regression analysis, linear models, and related methods. – Sage Publ.
- Fox, J. 2002. Linear mixed models: appendix to an R and S-PLUS companion to applied regression. – Sage Publ.
- Frair, J. L. et al. 2004. Removing GPS collar bias in habitat selection studies. – *J. Appl. Ecol.* 41: 201–212.
- Garshelis, D. L. 2000. Delusions in habitat evaluation: measuring use, selection, and importance. – In: Boitani, L. and Fuller, T. (eds), *Research techniques in animal ecology: controversies and consequences*. Columbia Univ. Press, pp. 111–163.
- Gillies, C. S. et al. 2006. Application of random effects to the study of resource selection by animals. – *J. Anim. Ecol.* 75: 887–898.
- Guisan, A. and Zimmermann, N. E. 2000. Predictive habitat distribution models in ecology. – *Ecol. Model.* 135: 147–186.
- Guisan, A. et al. 2002. Generalized linear and generalized additive models in studies of species distributions: setting the scene. – *Ecol. Model.* 157: 89–100.
- Hall, L. S. et al. 1997. The habitat concept and a plea for standard terminology. – *Wildl. Soc. Bull.* 25: 173–182.
- Hammond, P. S. et al. 1994. The diet of grey seals around Orkney and other island and mainland sites in northeastern Scotland. – *J. Appl. Ecol.* 31: 340–350.
- Hastie, T. J. and Tibshirani, R. J. 1990. Generalized additive models. – Chapman and Hall.
- Hjermann, D. O. 2000. Analyzing habitat selection in animals without well-defined home ranges. – *Ecology* 81: 1462–1468.
- Huffer, F. W. and Wu, H. L. 1998. Markov chain Monte Carlo for autologistic regression models with application to the distribution of plant species. – *Biometrics* 54: 509–524.
- Hutchinson, G. E. 1957. Concluding remarks. – *Cold Spring Harbor Symp. Quant. Biol.* 22: 415–427.
- Isaaks, E. H. and Srivastava, R. M. 1990. Applied geostatistics. – Oxford Univ. press.
- Johnson, C. J. et al. 2006. Resource selection functions based on use-availability data: theoretical motivation and evaluation methods. – *J. Wildl. Manage.* 70: 347–357.
- Johnson, D. H. 1980. The comparison of usage and availability measurements for evaluating resource preference. – *Ecology* 6: 65–71.
- Jolliffe, I. T. 1990. Principal component analysis: a beginner’s guide – I. Introduction and application. – *Weather* 45: 375–382.
- Jonsen, I. D. et al. 2003. Meta-analysis of animal movement using state-space models. – *Ecology* 84: 3055–3063.
- Keating, K. A. and Cherry, S. 2004. Use and interpretation of logistic regression in habitat selection studies. – *J. Wildl. Manage.* 68: 774–789.
- Keitt, T. H. et al. 2002. Accounting for spatial pattern when modelling organism-environment interactions. – *Ecography* 25: 616–625.
- Kernohan, B. J. et al. 1998. Use of an adaptive kernel home-range estimator in a GIS environment to calculate habitat use. – *J. Environ. Manage.* 53: 83–89.
- Latimer, A. M. et al. 2006. Building statistical models to analyze species distributions. – *Ecol. Appl.* 16: 33–50.
- Legendre, P. 1993. Spatial autocorrelation – trouble or new paradigm. – *Ecology* 74: 1659–1673.
- Lele, S. R. and Keim, J. L. 2006. Weighted distributions and estimation of resource selection probability functions. – *Ecology* 87: 3021–3028.
- Levin, S. A. 1992. The problem of pattern and scale in ecology. – *Ecology* 73: 1943–1967.
- Levins, R. 1966. The strategy of model building in population biology. – *Am. Sci.* 54: 421–431.
- Manly, B. F. J. et al. 2002. Resource selection by animals: statistical design and analysis for field studies, 2nd ed. – Chapman and Hall.
- Matthiopoulos, J. 2003a. Model-supervised kernel smoothing for the estimation of spatial usage. – *Oikos* 102: 367–377.
- Matthiopoulos, J. 2003b. The use of space by animals as a function of accessibility and preference. – *Ecol. Model.* 159: 239–268.

- Matthiopoulos, J. et al. 2004. Using satellite telemetry and aerial counts to estimate space use by grey seals around the British Isles. – *J. Appl. Ecol.* 41: 476–491.
- McConnell, B. J. et al. 1992. Satellite tracking of gray seals (*Halichoerus grypus*). – *J. Zool.* 226: 271–282.
- McConnell, B. J. et al. 1999. Movements and foraging areas of grey seals in the North Sea. – *J. Appl. Ecol.* 36: 573–590.
- McCullagh, P. and Nelder, J. A. 1989. *Generalized linear models*, 2nd ed. – Chapman and Hall.
- McLaughlin, J. F. and Roughgarden, J. 1992. Predation across spatial scales in heterogeneous environments. – Elsevier.
- Morales, J. M. et al. 2004. Extracting more out of relocation data: building movement models as mixtures of random walks. – *Ecology* 85: 2436–2445.
- Mysterud, A. and Ims, R. A. 1998. Functional responses in habitat use: availability influences relative use in trade-off situations. – *Ecology* 79: 1435–1441.
- Nielsen, S. E. et al. 2005. Modeling species occurrence and abundance; does probability of occurrence reflect population density? – *Ecography* 28: 197–208.
- Otter Research 2004. *Random effects in AD Model Builder. ADMB-RE user guide*. – Otter Research, Sidney BC.
- Overmars, K. P. et al. 2003. Spatial autocorrelation in multi-scale land use models. – *Ecol. Model.* 164: 257–270.
- Patterson, T. A. et al. in press. State-space models for individual animal movement. – *Trends Ecol. Evol.*
- Pinheiro, J. C. and Bates, D. M. 2000. *Mixed-effects models in S and S-plus*. – Springer.
- Prentice, R. L. and Pyke, R. 1979. Logistic disease incidence models and case-control studies. – *Biometrika* 66: 403–411.
- Priede, I. G. and Swift, S. M. 1992. *Wildlife telemetry; remote monitoring and tracking of animals*. – Ellis Horwood, New York.
- Ramsay, B. W. 2005. *Functional data analysis*. – Springer.
- Ripley, B. D. 1981. *Spatial statistics*. – Wiley.
- Rooney, S. M. et al. 1998. Autocorrelated data in telemetry studies: time to independence and the problem of behavioural effects. – *Mammal Rev.* 28: 89–98.
- Scott, J. M. et al. 2002. *Predicting species occurrences; issues of accuracy and scale*. – Island Press.
- Silverman, B. W. 1986. *Density estimation for statistics and data analysis*. – Chapman and Hall.
- Skaug, H. J. 2002. Automatic differentiation to facilitate maximum likelihood estimation in nonlinear random effects models. – *J. Comput. Graph. Stat.* 11: 458–470.
- SMRU 2005. *Scientific advice on matters related to the management of seal populations: 2005*. – SCOS, main advice report.
- Stein, A. and Corsten, L. C. A. 1991. Universal kriging and cokriging as a regression procedure. – *Biometrics* 47: 575–587.
- Stephenson, C. M. et al. 2006. Modelling establishment probabilities of an exotic plant, *Rhododendron ponticum*, invading a heterogeneous, woodland landscape using logistic regression with spatial autocorrelation. – *Ecol. Model.* 193: 747–758.
- Swihart, R. K. S. and Slade, N. A. 1985. Testing for independence of observations in animal movements. – *Ecology* 66: 1176–1184.
- Thomas, D. L. and Taylor, E. J. 2006. Study designs and tests for comparing resource use and availability II. – *J. Wildl. Manage.* 70: 404–412.
- Thompson, D. and Fedak, M. A. 2001. How long should a dive last? A simple model of foraging decisions by breath-hold divers in a patchy environment. – *Anim. Behav.* 61: 287–296.
- Tilman, D. and Kareiva, P. 1997. *Spatial ecology: the role of space in population dynamics and interspecific interactions*. – Princeton Univ. Press.
- Turchin, P. 1998. *Quantitative analysis of movement; measuring and modelling population redistribution in animals and plants*. – Sinauer.
- Vincent, C. et al. 2002. Assessment of Argos location accuracy from satellite tags deployed on captive gray seals. – *Mar. Mammal Sci.* 18: 156–166.
- Welsh, A. H. et al. 1996. Modelling the abundance of rare species: statistical models for counts with extra zeros. – *Ecol. Model.* 88: 297–308.
- White, G. C. and Garrott, R. A. 1990. *Analysis of wildlife radio-tracking data*. – Academic Press.
- Whittaker, R. H. L. S. A. and Root, R. B. 1973. Niche, habitat and ecotype. – *Am. Nat.* 107: 321–338.
- Wood, S. N. 2006. *Generalized additive models: an introduction with R*. – Chapman and Hall.
- Wright, P. J. et al. 2000. The influence of sediment type on the distribution of the lesser sandeel, *Ammodytes marinus*. – *J. Sea Res.* 44: 243–256.

Download the Supplementary material as file E5236 from
www.oikos.ekol.lu.se/appendix.