

The Use of Bayesian Model Averaging to Better Represent Uncertainty in Ecological Models

B. A. WINTLE,* †† M. A. McCARTHY,* † C. T. VOLINSKY,‡ AND R. P. KAVANAGH§

*School of Botany, University of Melbourne, Victoria, 3010, Australia

†Australian Research Centre for Urban Ecology, University of Melbourne, Victoria 3010, Australia

‡AT&T Labs Research, 180 Park Avenue, P.O. Box 971, Florham Park, NJ 07932-0000, U.S.A.

§Forest Research and Development Division, State Forests of New South Wales, P.O. Box 100, Beecroft, New South Wales 2119, Australia

Abstract: *In conservation biology, uncertainty about the choice of a statistical model is rarely considered. Model-selection uncertainty occurs whenever one model is chosen over plausible alternative models to represent understanding about a process and to make predictions about future observations. The standard approach to representing prediction uncertainty involves the calculation of prediction (or confidence) intervals that incorporate uncertainty about parameter estimates contingent on the choice of a “best” model chosen to represent truth. However, this approach to prediction based on statistical models tends to ignore model-selection uncertainty, resulting in overconfident predictions. Bayesian model averaging (BMA) has been promoted in a range of disciplines as a simple means of incorporating model-selection uncertainty into statistical inference and prediction. Bayesian model averaging also provides a formal framework for incorporating prior knowledge about the process being modeled. We provide an example of the application of BMA in modeling and predicting the spatial distribution of an arboreal marsupial in the Eden region of southeastern Australia. Other approaches to estimating prediction uncertainty are discussed.*

El Uso de Promedios de Modelo Bayesiano para Mejorar la Representación de la Incertidumbre en Modelos Ecológicos

Resumen: *La incertidumbre acerca de la elección de un modelo estadístico es raramente considerada en biología de la conservación. La incertidumbre de selección de modelo ocurre cuando se selecciona un modelo entre otros posibles modelos alternativos para representar el entendimiento de un proceso y para hacer predicciones acerca de observaciones futuras. El método estándar para representar la incertidumbre de predicciones implica el cálculo de intervalos de predicción (o confianza) que incorporan la incertidumbre en estimaciones de parámetros dependiendo de la selección de un modelo “mejor” seleccionado para representar la verdad. Sin embargo, este método de predicción basado en modelos estadísticos tiende a ignorar la incertidumbre de selección de modelo, lo que resulta en predicciones demasiado confiadas. El promedio de modelo Bayesiano (PMB) ha sido promovido en varias disciplinas como un medio simplificado para incorporar la incertidumbre de selección de modelo en la inferencia y predicción estadística. El PMB también proporciona un marco formal para la incorporación de conocimiento previo acerca del proceso a modelar. Proporcionamos un ejemplo de la aplicación de PMB en el modelado y predicciones la distribución espacial de un marsupial arbóreo en la región Edén del sureste de Australia. Se discuten otros métodos para estimar la incertidumbre de predicción.*

Model Uncertainty

A critical aspect of conservation biology involves predicting the efficacy or impacts of different management strategies. Biologists often seek to make predictions about un-

known quantities, such as the probability of occurrence or the future population size of a threatened species, based on a set of things that are known, such as aspects of the environment or the biology of a species. The true relationship between the attribute to be predicted and the

††email brendanw@unimelb.edu.au

Paper submitted December 11, 2001; revised manuscript accepted May 28, 2003.

information at hand is seldom known, so a model is relied upon to formalize a judgment about the relationship. A model will generally comprise two parts: a structure and a set of parameter estimates specific to that structure. In the case of population models, model structure may be defined by the presence or absence of stochasticity and independent variables and functions that describe features such as competition, density dependence, spatial structure, and genetic drift. In the case of habitat models, the choices of link function, variables, transformations, and interactions define the structure. For example, Lindenmayer (1990a) developed a habitat model to predict the probability of occurrence (y) of the greater glider (*Petauroides volans*) as a function of the age of the forest (a) and the number of trees (>0.5 m in diameter) with cavities in a given 3-ha cell (n):

$$\ln\left(\frac{y}{1-y}\right) = -0.993 + 1.106a + 0.554n. \quad (1)$$

In this example, the structural choices or assumptions made by the authors include (1) the choice of the variables a and n and the exclusion of a range of other possible predictors relating to demographic processes or the foraging, denning, and metabolic requirements of the species, (2) a logit link function rather than probit or other alternatives, and (3) simple linear relationships involving no transformation of or interaction between the predictor variables.

Clearly, a large range of alternative models might result from the combination of the possible assumptions listed above. In practice, the structural features of the model are often determined by some data-driven search for the "best" model from the range of possible models. Once the model structure is chosen, the estimation of parameters (such as fitting of regression coefficients) follows.

The search for the best model structure recognizes the existence of more than one candidate model. This implies a level of uncertainty associated with the choice of model structure that is usually ignored when it comes to making predictions and computing prediction intervals with the chosen best model. The best model is assumed to have correct structure, and uncertainty in parameter estimates is the only form of uncertainty considered in making predictions. This assumption is tenuous given that errors arising from uncertainty about the structure of the model are likely to be far worse than those arising from other sources (Chatfield 1995). This is particularly relevant in ecological models, in which uncertainty about model structure is usually high (Conroy et al. 1995). Thus, the assumption that a single "true" model exists is rarely justified (Reichert & Omlin 1997), and alternative model structures can result in very different predictions (Pascual et al. 1997; Beissinger & Westphal 1998).

In some instances, the choice of model structure may be obvious, based on the subject matter or the design of a

study. For example, geometry or allometric relationships may determine model structure in physiological models (Porter et al. 2000). This is rarely the case in observational studies in ecology, however, where there is often a great deal of ambiguity about the processes involved in determining the true value of the response. Where there is ambiguity about the true model structure in regression analysis, automated variable-selection methods are commonly used to search large sets of candidate predictors for plausible model structures (e.g., Loyn et al. 2001). The propensity of automated variable-selection methods to choose variables with no real relationship to the dependent variable (Derksen & Keselman 1992) is an example of a common source of structural uncertainty that is seldom acknowledged.

If a number of possible model structures fit the observed data almost as well as the chosen best structure, there will be considerable uncertainty about which model is in fact the best. If a plausible alternative model structure results in predictions that are different from those of the chosen best model, there is some risk involved in bravely ignoring the alternative argument(s) and putting all eggs in one basket. A more conservative approach is to consider any given model as a sample from a large set of competing models and then use data to weight or assign a degree of belief to the competing models.

Broadly speaking, approaches for incorporating model-selection uncertainty in prediction fall into three categories. The first approach is to do nothing. This is the simplest and most popular method, based on a data-driven search for a single best model and selection criteria that appeal to the conceptual trade-off between bias and variance (Burnham & Anderson 1998). Prediction and estimation of precision are undertaken on the assumption that the best model is the true model and uncertainty exists solely in the estimation of model parameters (or coefficients in the regression context).

A second approach is discrete model averaging (Draper 1995). The fundamental goal of model averaging is to arrive at a set of plausible models that can be weighted according to some criteria, usually associated with their likelihood given a set of data. This weighting reflects the degree to which each model is trusted. Predictions are then derived as a weighted average of the predictions from each model in the plausible set, and prediction uncertainty is estimated as a weighted sum of the within- and between-model variance. Discrete model averaging is "discrete" because weighted averages are constructed for a predetermined subset of all possible models. In the example below, we used a method to define a discrete sample of models from the full model space.

A third approach is continuous model averaging. Draper (1995) introduced the concept of "model expansion," whereby a single structural choice is expanded in directions suggested by the context or the analytic search that specified the favored model. Markov chain

Monte Carlo (MCMC) methods provide a vehicle for “continuous” (rather than discrete) expansion of the model structure by iteratively sampling from the full range of prior distributions of parameters in the construction of posterior parameter distributions. Applying MCMC in regression problems that involve numerous candidate predictors, possible interactions, and transformations can be complex, and ensuring convergence and adequate mixing may be problematic. This complexity is compounded by a lack of access to MCMC applications in standard statistical software. This approach warrants further investigation for ecological application, however, because some suitable packages are being developed, but this is outside the scope of our paper. See Fleishman et al. (2001) for one conservation application. Detailed descriptions of the use of MCMC in model-selection problems are available (George & McCulloch 1995; Raftery et al. 1997).

A Framework for Dealing with Model Uncertainty

The frequentist approach to inference and prediction does not adapt naturally to coping with model uncertainty (Chatfield 1995), owing to an absence of formal methods for deriving and incorporating posterior model probabilities. Because the bulk of biologists tend to favor frequentist methods, this may partly explain why model uncertainty is so seldom reported in biological examples of inference and prediction. Buckland et al. (1997) describe an ad hoc approach based on the bootstrap to account for uncertainty in model selection in a frequentist way. We show, however, that model averaging fits more naturally within the formal Bayesian theoretical framework, making it an appealing approach to incorporating model-selection uncertainty in inference and prediction. In this section we aim to illustrate the simplicity with which estimating model uncertainty may be handled in a Bayesian way. Ellison (1996) and Wade (2000) review the general use of Bayesian methods in conservation biology.

A model space (M) may be defined as a set of possible model structures (S_i) and a set of parameter vectors (θ_i) that are specific to each structure. After collecting some data (D), a biologist might attempt to predict the probability of a future outcome (Δ), such as the occurrence of a species ($\Delta = 1$), via the probability model (Raftery 1996):

$$\Pr(\Delta | D) = \sum_{i=1}^I \Pr(\Delta | S_i, D) \Pr(S_i | D), \quad (2)$$

where S_i are the competing models from the model space M ; $\Pr(S_i | D)$ is the posterior model probability and represents the degree of belief in model S_i given the data collected and the prior belief (if there was one) before the data were collected; and $\Pr(\Delta | S_i, D)$ represents a

(posterior) prediction of Δ according to model S_i and the data (D) used to parameterize the model.

The $\Pr(\Delta | D)$ represents a prediction for Δ that is a weighted average of the predictions from each of the models S_i . The weight assigned to the prediction from any one of the S_i is the posterior probability (or degree of belief) $\Pr(S_i | D)$ in that model, given the data and prior knowledge. It is the incorporation of the posterior model probability in the prediction of Δ that provides the explicit recognition of model uncertainty that is missing in classical frequentist analysis. The main questions are, then, how are the competing models (S_i) identified, and how should posterior probabilities be assigned?

Identifying Competing Models

In some instances, a small set of competing models may represent well-defined schools of thought or competing hypotheses about a biological process. For example, when the long-term viability of a species is analyzed by PVA, there may be a limited number of theoretical models to explain the population dynamics of the species being considered. In these “discrete” situations, the only impediment to a simple BMA analysis is the identification of a coherent scheme for assigning posterior model probabilities to candidate models. In many modeling situations, however, the number of candidate variables, functional forms, and transformations can be great, in which case the number of possible model structures becomes enormous. Approaches have been developed to reduce the size of the problem. The “leaps and bounds” algorithm (Furnival & Wilson 1974), used in the example below, provides a method for identifying the sample of models representing the full model space. Selection of stochastic search variables (George & McCulloch 1995), MCMC, and importance sampling are alternative approaches to sampling from the model space, although we could not find any applications of these three approaches in logistic regression.

Assigning Model Posterior Probabilities

Bayes's Theorem is easily extended to enable the calculation of the posterior probability of model S_i (Hoeting et al. 1999):

$$\Pr(S_i | D) = \frac{\Pr(D | S_i) \Pr(S_i)}{\sum_{k=1}^S \Pr(D | S_k) \Pr(S_k)}, \quad (3)$$

where $\Pr(S_i)$ is the prior probability (or belief) in model S_i , and

$$\Pr(D | S_i) = \int \Pr(D | \theta_i, S_i) \Pr(\theta_i | S_i) d\theta_i \quad (4)$$

is the integrated likelihood of model S_i , which provides a measure of how much the data support each model. The θ_i is the vector of parameters of model S_i (e.g., for

regression $\theta = [\beta, \sigma^2]$; $\Pr(\theta_i | S_i)$ is the prior density of θ_i under S_i ; and $\Pr(D | \theta_i, S_i)$ is the likelihood of observing the data given the estimated parameters specific to S_i . Although this appears to be an imposing number of parameters, likelihoods, and conditional distributions to deal with, it can be seen from the regression example below that the necessary elements are often able to be directly estimated by or are by-products of model-fitting in standard statistical packages.

The integral in Eq. 4 can be hard to compute (Hoeting et al. 1999). A simpler approach is to use the “BIC (Bayesian information criterion; Schwarz 1978) approximation” based on the Laplace method (Tierney & Kadane 1986). The BIC approximation for determining posterior model probabilities is accurate when sample sizes are large enough (>20 times the number of predictors) (Kass & Raftery 1995), and its use has been justified by several authors (e.g., Hoeting et al. 1999; Volinsky & Raftery 2000). Therefore, we utilized the BIC approximation as a simple means of assigning posterior model probabilities.

Specifying Prior Distributions

Bayesian model averaging requires the specification of prior distributions for the parameters $\Pr(\theta_i | S_i)$ and the models $\Pr(S_i)$. The use of prior knowledge or beliefs has been interpreted as both the strength and weakness of the Bayesian approach to inference and prediction. On the one hand, the formal role of prior knowledge in Bayesian analysis provides a natural framework for combining new and existing information in an adaptive management cycle (Ellison 1996). However, prior distributions can be difficult to specify (Clyde 1999) and, if handled poorly, may lead to undesirable behavior of the posterior distribution (Berger 1985). Specification of prior distributions is often subjective (Ellison 1996), although it is clear from Bayes's theorem that if the evidence provided by the data is compelling, the likelihood will dominate the posterior and the prior will have little effect. If the data are not compelling, there are no grounds on which to alter current (possibly subjective) beliefs, so it makes sense that the prior should dominate the posterior. The advantage of the Bayesian approach to model averaging is that it provides a coherent and intuitive theoretical framework, encourages the quantification of prior knowledge, and, in doing so, forces transparency about subjective assumptions. It also encourages users to identify forms of objective prior information such as alternative data sources. The disadvantage is that it requires the specification of prior distributions, even when no strong prior information exists. In situations such as exploratory data analysis, where little if any prior belief exists, it is necessary to specify objective (or uninformative) prior distributions in order to use the Bayesian formulation without biasing results. Considerable work has been invested in developing objective

prior distributions for Bayesian analysis (Jeffreys 1961; Raftery 1996; Clyde 2000).

When prior information about the importance of a variable is available for model structures with a coefficient associated with each predictor (e.g., generalized linear models), a prior probability on model S_i can be specified as (Hoeting et al. 1999):

$$\Pr(S_i) = \prod_{j=1}^p \pi_j^{\delta_{ij}} (1 - \pi_j)^{1 - \delta_{ij}}, \quad (5)$$

where π_j is the prior probability that variable j has a non-zero coefficient ($\theta_{ij} \neq 0$), and δ_{ij} (taking a value of 0 or 1) indicates whether variable j is included in model S_i . Assigning $\pi_j = 0.5$ for all j corresponds to a uniform prior across the model space.

Using a model-selection criteria such as AIC (Akaike's information criteria; Akaike 1973) or BIC precludes the need to explicitly specify a prior distribution for the parameters $\Pr(\theta_i | S_i)$. These criteria correspond to prior distributions that are proper but uninformative (Clyde 2000). For example, BIC defines an implicit prior distribution—a normal distribution centered around θ_i —with the amount of information in the prior equal to the average amount of information in one observation (Volinsky & Raftery 2000).

Debate over the use of AIC and BIC as the optimal model-selection criteria is unresolved, although it is clear that BIC favors smaller, less complicated model structures. The use of BMA based on BIC has precedent in the statistical literature (Raftery 1996; Hoeting et al. 1999), based on the fact that BIC is an approximation to twice the log of the Bayes factor, the standard test statistic in Bayesian hypothesis testing. The usual justification of AIC provides no way of taking into account model uncertainty (Clyde 2000). Clyde (2000) provides a convenient representation of model-selection criteria in which a change in a single parameter allows the user to implement AIC, BIC, or several other options. Clyde's formulation (CIC: calibrated information criterion) provides a Bayesian justification for the entire family of criteria, derived from Jeffrey's prior distribution (Jeffreys 1961).

In the example below, we used objective prior distributions for $\Pr(\theta_i | S_i)$ and $\Pr(S_i)$ developed by Clyde (2000) that make use of readily available summary statistics such as AIC and BIC, derived from standard statistical packages.

An Example of Bayesian Model Averaging: Inferring Habitat Preferences and Predicting the Spatial Distribution of the Greater Glider

Background

Wildlife habitat models that quantify the relationship between the occurrence of species and habitat

characteristics have been used throughout the world in resource analysis and conservation planning exercises (Pereira & Itami 1991; Mladenoff & Sickley 1998; Queensland Government 1998; Li et al. 1999; Loyn et al. 2001). The purpose of our example is to demonstrate how the single best model approach to habitat modeling can be extended easily to the BMA approach, which enables incorporation of model-selection uncertainty in inference, predictions, and prediction intervals. We demonstrate how the structural uncertainty of models affects their inference about key species-habitat relationships, their predictive accuracy, and the coverage of their prediction intervals.

The greater glider (*Petauroides volans*; Kerr 1792) is a marsupial glider endemic to eastern Australia. Being an obligate cavity nester, it is dependent on elements of old forest (Lindenmayer et al. 1990a) and consequently is of conservation concern in areas subject to timber harvesting. Habitat models were constructed using logistic regression (McCullagh & Nelder 1989) to describe the relationship between the probability of occurrence of the greater glider and habitat attributes such as vegetation type, forest age, climate, and topography. There are numerous precedents for using logistic regression and other generalized linear models (GLMs) in habitat analyses (e.g., Nicholls 1989; Pereira & Itami 1991; Buckland & Elston 1993). However, no example of an explicit treatment of model-selection uncertainty in wildlife habitat modeling could be found, despite the compelling arguments and methods proposed by Buckland et al. (1997).

The Species

The greater glider is an arboreal, folivorous, gliding marsupial whose distribution extends from southern Australia to just north of the tropic of Capricorn (McKay 1995). A considerable amount of research has been devoted to exploring its habitat preferences (e.g., Kavanagh & Lambert 1990; Lindenmayer et al. 1990a), distribution (New South Wales National Parks and Wildlife Service 1998), population ecology (Possingham et al. 1994), physiology (Foley et al. 1990), sociology (Henry 1984), and response to logging (Kavanagh 2000).

The Habitat Variables

Cork and Catling (1996) categorize the bulk of studies on arboreal mammals into those pursuing the hypothesis that the foliar nutrient status of the forest is the prime determinant of habitat quality (Braithwaite 1983) and those that put equal or greater emphasis on variables related to structural characteristics of the forest. Significant relationships with variables describing forest structural attributes such as stand basal area, presence of trees with hollows, old-growth forest patch size, and time since logging have been reported (e.g., Lindenmayer et al. 1990a; Kavanagh & Bamkin 1995). Other researchers have found

significant relationships between greater glider presence and forest type, tree species, and foliar nutrient indices (e.g., Braithwaite et al. 1988; Kavanagh & Lambert 1990). Lindenmayer et al. (1990b) and Pausas et al. (1995) found that indirect surrogates of shelter, nutrition, or mobility, such as slope and topographic position, are good predictors of arboreal marsupial habitat. Climatic variables are likely to affect the distribution of greater glider habitat on a broad scale because of thermal and physiological constraints (Porter et al. 2000). We compiled a list of uncorrelated ($R \leq 0.6$) candidate predictors based on these habitat analyses (Table 1).

Study Area and Glider Surveys

Fauna surveys were conducted in southeastern Australia (Fig. 1) in 1992 and 1994 by Kavanagh and Bamkin (1995) and Kavanagh (1997), providing 219 observations of greater glider presence or absence for model building and 187 separate and independent observations of presence or absence for model testing.

Model Building

PRIOR DISTRIBUTIONS

In many data-exploration exercises, including our example, there may be little information available for eliciting prior distributions $\Pr(\theta_j | S_i)$ and $\Pr(S_i)$. Moreover, where the number of predictor variables is large (> 10) and transformations and polynomial terms are considered likely, the elicitation of subjective prior distributions can become extremely complicated. In this context, “uninformative” prior distributions that, a priori, make all models and parameters equally likely are appealing.

As discussed previously, using model-selection criteria such as AIC and BIC removes the need to explicitly specify $\Pr(\theta_{ij} | S_i)$. We implemented BMA with both AIC and BIC with Clyde’s CIC objective prior formulation, amounting to the specification of a uniform prior $\Pr(S_i)$ and prior distributions with uninformative proper parameters.

POSTERIOR DISTRIBUTIONS

In our regression problem, uncertainty about the predicted probability of occurrence of a greater glider amounts to uncertainty about estimates of the coefficients (θ_j) in a single model and the additional uncertainty about which model to use. Consequently, the posterior distributions of interest are (Clyde 2000)

$$\Pr(\theta_j | S_i, D) = N(\hat{\theta}_j, I(\hat{\theta}_j)^{-1}), \quad (6)$$

and

$$\Pr(S_i | D) = \frac{\exp\{1/2(\text{Dev}_i - d_i \log(c))\}}{\sum_{i=1}^I \exp\{1/2(\text{Dev}_i - d_i \log(c))\}}, \quad (7)$$

Table 1. Map layers used as candidate predictors of greater glider habitat.

Variable name (abbreviation)	Description
Foliar nutrient index (foliar.nut)	binary variable derived from Braithwaite's foliar nutrient index (New South Wales National Parks and Wildlife Service 1998), which is based on vegetation mapping (Keith & Bedward 1999); 10 categories reduced to 2 by imposing a threshold at category 4 (on the basis of preliminary data analysis)
Old-growth and mature forest in a 150-m radius* (old150)	context* variable that indicates the number of 25-m grid cells classified as old growth within a 150-m radius of the target cell
Topographic position (topo.pos)	measure of the position of a cell on a continuum between 0 (gully) to 100 (ridge) (New South Wales National Parks and Wildlife Service 1998)
Wet forest within 150 m* (wetforest)	measure of the proportion of forest in a 150-m radius classified as wet forest (Keith & Bedward 1999)
Rainforest within 150 m* (rf150)	measure of the proportion of forest in a 150-m radius classified as rainforest (Keith & Bedward 1999)
Dry forest within 150 m* (dryforest)	measure of the proportion of forest within 150 m classified as dry sclerophyll forest (Keith & Bedward 1999)
Mean annual temperature (mean.temp)	derived from a DEM (Land Information Centre 1997) using the ESOCIM module of ANUCLIM (Hutchinson et al. 1999)
Wetness index (wetness)	derived from the DEM (Land Information Centre 1997) to indicate the volume of water draining to a 100-m cell and its ability to retain the water based on local slope
Solar radiation index (solar)	derived from the DEM (Land Information Centre 1997) and ESOCIM (Hutchinson et al. 1999) to indicate the amount of solar radiation that falls on each grid cell by allowing for shade and shadow due to terrain and atmospheric scattering

*Context variables are those in which the value of each cell is a function of the values of surrounding cells in a specified radius.

where Dev_i is the model deviance ($-2 \times [\log \text{likelihood}]$) under the null model minus the deviance under model S_i ; d_i is the number of estimated parameters of model S_i ; $I(\hat{\theta}_i)$ is the observed Fisher information for S_i evaluated at the maximum-likelihood estimates (MLEs) of $\hat{\theta}_i$; and $\log(c)$ is a calibration constant determined by the model-selection criteria used (Clyde 2000). For BIC, $\log(c) = \log(n)$, and for AIC, $\log(c) = 2$. Fisher information $I(\hat{\theta}_i)$ for model S_i is $\mathbf{X}'_i \mathbf{V}(\hat{\theta}_i) \mathbf{X}_i$, where $\mathbf{V}(\hat{\theta}_i)$ is the covariance matrix for D , which can be calculated in any statistical package that fits GLMs.




Figure 1. Glider survey locations in southeastern Australia.

In our example, we sampled from the posterior distributions of coefficients and models to examine perceived species-habitat relationships, including the relationship between glider occupancy and the proportion of old and mature forest within the surrounding area. We compared posterior distributions obtained under AIC and BIC prior distributions to illustrate the combined effect of model-selection criteria and parameter and model uncertainty on model inferences about species-habitat relationships.

IMPLEMENTATION

Despite the rather arduous theoretical development and justifications of BMA, the implementation in the regression context is simple. We implemented BMA for a logistic-regression analysis of greater glider habitat relationships with SPLUS (MathSoft 1997) code available from the BMA website: <http://www.research.att.com/~volinsky/bma.html> entitled BMA.GLM. The BMA.GLM function incorporates Clyde's CIC prior and posterior model distributions, allowing BMA under both BIC and AIC prior distributions. The code utilizes the leaps-and-bounds algorithm (Furnival & Wilson 1974) to rapidly identify a suite of models representing the full model space. The BMA.GLM function enables the selection of polynomial terms in competing models for fitting non-linear relationships. The BMA.GLM function also allows the explicit allocation of prior probabilities that particular candidate variables are non-zero, where there are grounds to assign them. It requires minimal running time and normal SPLUS model-fitting inputs. The BMA.GLM function reports a subset of candidate models, provides estimates of posterior model probabilities, MLEs of coefficients, and

standard errors of MLEs specific to each model, based on both AIC and BIC model-selection criteria. Inferences and predictions using the two criteria may then be compared. Other code is available at the BMA website to implement BMA for other linear models and survival analysis. These methods are discussed in more detail by Hoeting et al. (1999).

Running BMA.GLM requires only that the user specify a vector of n binary observations (y) that indicate the presence or absence of the species at the n survey locations, an $n \times p$ matrix (X) corresponding to the measured val-

ues of the p predictors at each of the n survey locations. The number of predictors that will take polynomial forms, the GLM family, and the model-selection criteria to be used (AIC or BIC) are all specified as arguments in the call to BMA.GLM from within SPLUS.

Model Uncertainty and Inference about Species-Habitat Relationships

We identified 62 competing models under AIC and 42 under BIC with the leaps-and-bounds component of

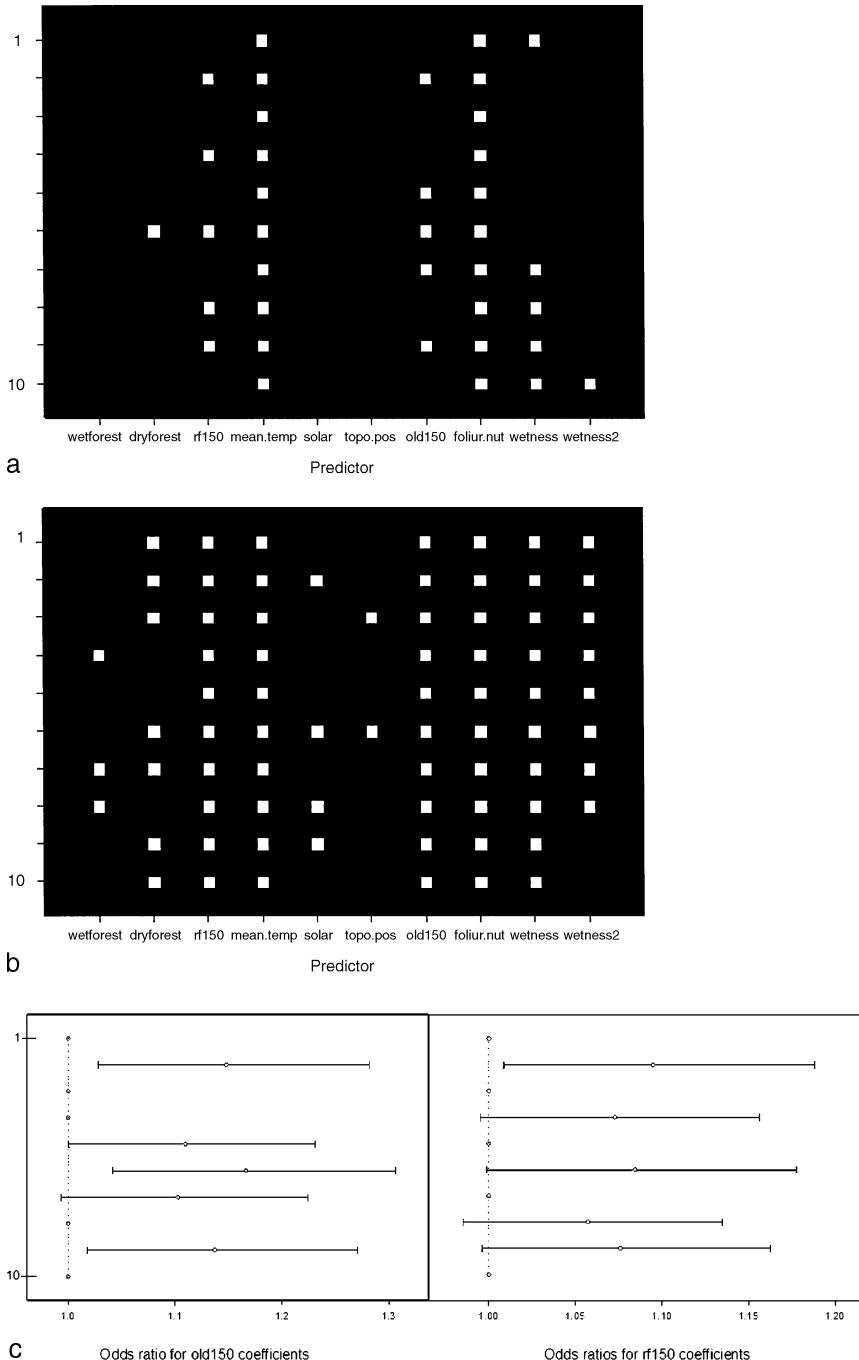


Figure 2. Plots of model space for (a) Bayesian information criterion (BIC) and (b) Akaike's information criteria and (c) corresponding maximum-likelihood estimates (MLEs) of odds ratios for two habitat attributes under BIC. The inclusion of particular predictors (listed on the x-axis) in any given model (models 1-10 are listed on the y-axis) is indicated by a white square. Models 1-10 (on the y-axis) are presented in rank order of posterior probability (that the particular model is the "true" model). Plot (c) demonstrates the uncertainty in model-specific estimates of coefficients for old150 and rf150 by "linking" the BIC model space to model-specific MLEs of odds ratios for the two habitat predictors. Horizontal bars represent 95% confidence intervals on MLEs of odd ratios. An odds ratio of 1 (represented by the dashed vertical line) implies that variable has no effect on species occupancy. Predictor abbreviations are defined in Table 1.

BMA.GLM. Ten of the most likely models from the two model spaces (defined under BIC and AIC prior distributions) are presented in Fig. 2. The posterior probabilities of the top 10 models were separated by <0.07 under both selection criteria, indicating a high degree of model uncertainty. The extra leniency afforded by the AIC approach to inclusion of variables resulted in the selection of larger, more complex models (Fig. 2). A number of the variables consistently present in AIC-selected models are only occasionally selected in BIC models. This highlights the magnitude of model uncertainty that may go unnoticed when only one model-selection criteria is used to construct models, as is often the case in species-habitat analyses.

Inference under Model Uncertainty

For compactness, we focused on the effect of model uncertainty on inferences about two candidate predictors of greater glider habitat: the proportion of old and mature forest in the surrounding area (old150) and the proportion of rainforest in the surrounding area (rf150). To avoid cumbersome graphical representations of our results, we focus this analysis on the top 10 models selected under both AIC and BIC model-selection approaches, though under normal circumstances a greater number of plausible models could easily be considered. Models selected by AIC contained old150 and rf150 much more often than those selected by BIC (cf. Fig. 2a & 2b). Under AIC both variables were always included in the 10 most likely models, whereas under BIC there was a great deal of uncertainty about the importance of the two variables, with 50% representation of both variables in the top 10 models. Unlike other predictors such as mean temperature, which was consistently represented in both the AIC and BIC models, there was more uncertainty about the importance of the two predictors old150 and rf150 than the AIC approach would indicate in isolation. Between-model variation in the (non-zero) MLEs of old150 and rf150 coefficients under BIC was substantial (Fig. 2c). Non-zero MLEs of odds ratios for the effect of old150 on the probability of observing a greater glider varied between 1.1 and 1.18. The odds ratio for a given variable in a given model represents the predicted multiplicative increase in the probability of a positive outcome (e.g., occurrence of a greater glider) resulting from a unit increase in the independent variable (Agresti 1996). For example, an odds ratio of 1.15 (for old150 in the second-best BIC model) indicated that a one-unit (10%) increase in the proportion of old forest in the neighborhood of a site results in a 15% increase in the probability of a greater glider occurrence at the site. The horizontal lines correspond to a 95% confidence interval on the MLE of the odds ratio.

The magnitude of the coefficient uncertainty varied between models, as did model-specific inferences about the effect of old150 on the odds of finding a greater glider un-

der the BIC-selected models (Fig. 3a). Uncertainty about model-specific inference was less apparent in the AIC-selected models (Fig. 3b). The middle layer of the plots in Fig. 3 represents model uncertainty and was obtained by jointly sampling the model posterior distribution (conditional on the inclusion of old150) and point estimates of coefficients (ignoring coefficient uncertainty). These plots show uncertainty about the coefficient estimate for old150 between models that contained it. Obviously, the level of model uncertainty associated with the coefficient would increase as a result of sampling from models that did not contain it (i.e., where $MLE_{old150} = 0$).

Finally, combined model and parameter uncertainty was demonstrated by the density of 10,000 samples taken jointly from $Pr(S_i | D)$ and $Pr(\theta_i | D, S_i)$, conditional on the inclusion of old150. This density was equivalent to the model-averaged posterior distribution of the old150 odds ratio. A notable feature of the BIC model-averaged distribution of old150 was that the mass of the distribution shifted left, to better represent the combined inference about old150 across models that include it (Fig. 3a). It is also notable that the absolute magnitude of variation in the BMA estimate of the old150 odds ratio was not substantially greater than the variation in the single, best estimate under both AIC and BIC. This is partly because the standard errors of MLEs for the old150 parameter were lower in some of the plausible alternative models in the model space and because models that did not include old150 were not incorporated in model-averaged inference. Finally, it is apparent that inference based on AIC alone would involve less hedging against model uncertainty than inference based on BIC or both selection criteria.

Comparing the Predictions of the Best Model and Bayesian Model Averaging

Predictions for independent model-testing sites were made with the best models obtained through AIC- and BIC-based models selection and by using BMA under both AIC and BIC prior distributions. The BMA predictions were made by obtaining 10,000 samples jointly from posterior distributions $Pr(S_i | D)$ and $Pr(\theta_i | D, S_i)$. Each sample represents a vector of coefficient values that may be applied in a logistic-regression model to make a prediction about the probability of greater glider occurrence at each of the 187 model testing sites. A Bayesian predictive distribution of 1000 predictions was constructed at each of the 187 test sites, from which a mean prediction and 90% probability interval were derived.

Discrimination is a measure of the ability of model predictions to distinguish correctly between presence and absence sites. The area under the receiver operating characteristics (ROC) curve has become a widely accepted measure of binary model discrimination (Fielding & Bell 1997). The area under the ROC curve can range between

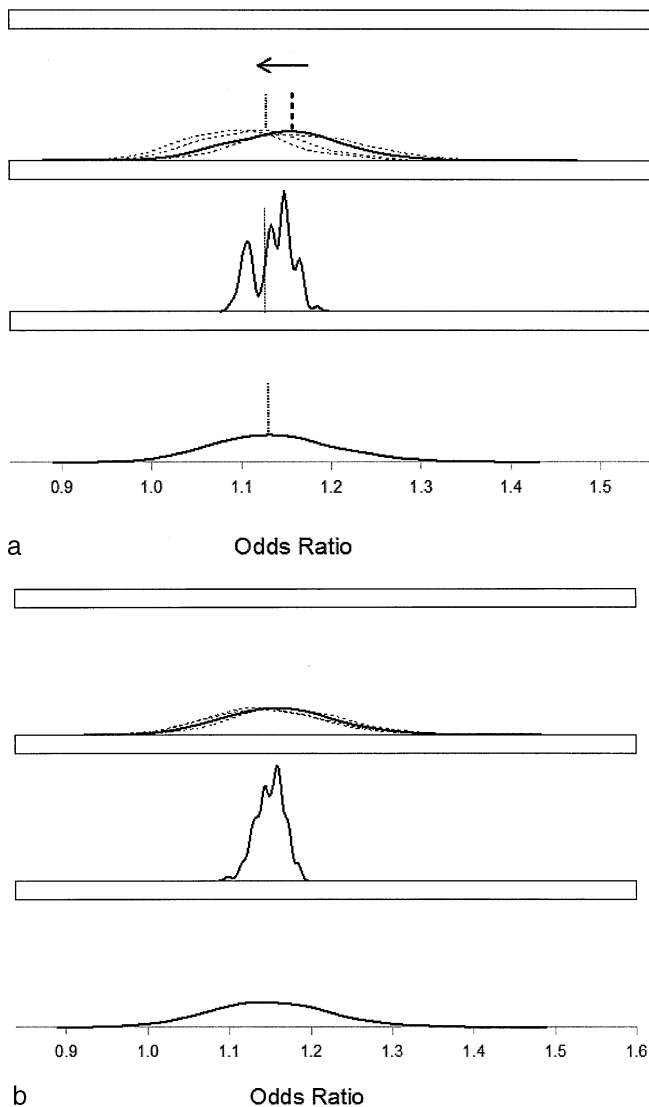


Figure 3. Distributions of odds ratios for *old150* incorporating parameter, model, and combined model-parameter uncertainty (conditional on the inclusion of *old150*) under both (a) Bayesian information criterion and (b) Akaike's information criteria. Parameter uncertainty (the top layer) is represented by the distributions of samples from the model-specific coefficient posterior distributions $Pr(\theta_i | D, S_i)$ for the four best models (the posterior distribution of the single best model is represented by a solid black line). Model uncertainty (the middle layer) is represented by the posterior distribution $Pr(S_i | D)$. Combined parameter-model uncertainty about the odds ratio for *old150* (the bottom layer) is the posterior distribution $Pr(\theta_i | D)$, derived by jointly sampling $Pr(\theta_i | D, S_i)$ and $Pr(S_i | D)$ 10,000 times. The lines drawn vertically at the modes of the best-model (dashed) and BMA (dotted) posterior distributions in (a) highlights the shift in inference about the odds ratio of *old150* from the best model to the Bayesian model averaging model.

0 and 1, where a score of 1 implies perfect discrimination and a score of 0.5 or less implies predictive discrimination that is no better than a random guess. Predictions derived from the best-model and BMA approaches were tested against observations at the 187 test sites, with the area under the ROC curve as an index of predictive performance. The best-model and model-averaged predictions based on both AIC and BIC prior distributions performed reasonably well with the area under the ROC curve varying between 0.78 and 0.80 for all four sets of predictions.

Comparing Prediction Intervals Derived from Bayesian Model Averaging with Those of the Best Model

As distinct from measures that indicate the accuracy of predictions (such as the area under the ROC curve), predictive coverage describes the optimism of confidence intervals. By comparing model prediction intervals, it is possible to discern how overconfident we were about the accuracy of our predictions.

Prediction intervals for the predictions of any single model are easily computed with the point-wise standard errors returned in a standard statistical package (Agresti 1996). Ninety percent prediction intervals based on the best AIC and BIC model predictions were computed for all 187 model testing sites. Computation of prediction intervals for a BMA prediction requires the integration of the intervals for each of the competing models. Prediction intervals for BMA predictions were based on the 1000 predictions at each model testing location described above. Choosing upper and lower limits such that 5% of the predictive density (50 predictions) lay above the upper limit and 5% lay below the lower limit identified a 90% probability interval for each prediction.

Predictive coverage measures the performance of prediction intervals according to how often the intervals contain the new observations. Measuring the predictive coverage of intervals that bound the probability of a binary outcome is complicated by the fact that observations can only be either 0 or 1, but predictions are necessarily probabilistic (0-1). Consequently, observations were grouped into decile ranges of prediction to test how well the observed proportion of positive observations concurred with the confidence intervals for the predicted proportion (Fig. 4). This approach to testing binary model predictions is the foundation of the Hosmer-Lemeshow test and the calibration plot, two commonly used methods for testing how well model predictions are calibrated (Miller & Hui 1991; Pearce & Ferrier 2000). For BIC models, the BMA prediction interval is substantially broader than the interval for the single most likely model selected under BIC because of the large amount of uncertainty in BIC models (Fig. 4). Bounds on AIC and BMA predictions (not presented) were slightly narrower, though still broader than the best-model prediction intervals. The observed proportions of greater glider presences fell within

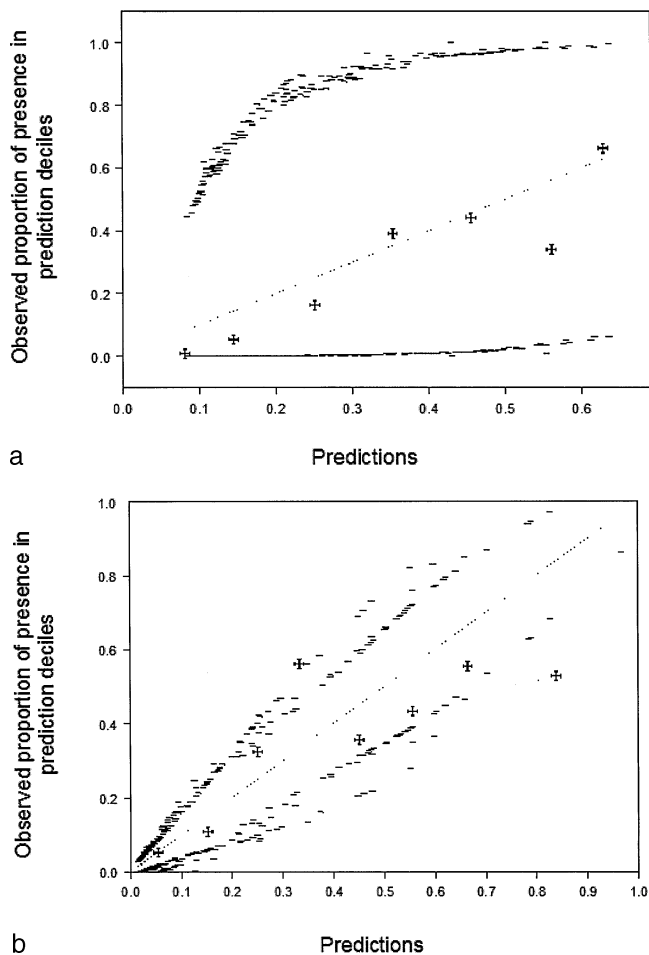


Figure 4. Ninety percent probability intervals for (a) Bayesian model averaging (BMA) and (b) best-model predictions based on Bayesian information criterion (BIC) model selection. Ninety percent prediction intervals are presented for each of the 187 predictions made for model testing sites and are defined by horizontal dashes. Ninety percent prediction intervals for the BMA predictions are based on all 42 plausible models identified under BIC model selection using the the SPLUS function BMA.GLM. Crosses are the observed proportion of greater glider presences in groups defined by prediction deciles. A perfect match between predictions and proportions of greater gliders observed in prediction deciles would see the crosses lying on the dotted line that runs diagonally through each plot.

the BMA prediction intervals in all instances, whereas the best-model interval bounded the observed proportion less often. This result is consistent with those of Raftery (1996), Raftery et al. (1997), and Hoeting et al. (1999).

Interpretation

Our example demonstrates that inference about species-habitat relationships based on one model alone is likely

to be naïve about the real uncertainty inherent in the analysis. Based on the results of our example, if the ramifications of inferring the dependence of greater gliders on old forest were large, ignoring model uncertainty about such inference would be perilous. Given the ease with which BMA.GLM is applied and the degree of insight it may provide into model uncertainty, it is clearly a useful tool for assessing species-habitat relationships and uncertainty about the nature of the relationships.

In contrast to the studies by Draper (1995), Hoeting et al. (1999), and others, the point-wise predictions of the best model showed predictive performance in terms of discrimination equivalent to those of the model-averaged predictions. This may be due in part to the dominance of one or two predictors in all of the models or to lack of complete independence between the model-building and model-testing data.

The usefulness of predictions can be thought of both in terms of the proportion of times that the mean prediction is correct and how well uncertainty about future observations is estimated. In our example, the BMA prediction intervals exhibited better coverage than those derived from the best models. The inadequacy of standard prediction intervals has been documented in a number of examples based on real data (Draper 1995; Raftery et al. 1997; Hoeting et al. 1999) and simulated data (Hurvich & Tsai 1990; Raftery et al. 1997).

Despite showing better predictive coverage than the single-best-model prediction intervals, the BMA intervals did not account for all forms of uncertainty inherent in models and model predictions. The assumption that there is no error in mapped variables and greater glider observation data likely contributed substantially to an underestimation of prediction uncertainty (Elith et al. 2002). Methods for incorporating these types of uncertainty into prediction intervals are poorly developed in ecology, despite some precedents in the medical sciences (Magder & Hughes 1997).

Conclusions

This work represents the first attempt to explicitly estimate and incorporate model structural uncertainty into the inference and predictions of species-habitat models. Although we did not demonstrate that BMA predictions are necessarily more accurate, we did demonstrate that BMA provides a simple means for achieving a more satisfactory treatment of inferential and prediction uncertainty than currently accepted methods. On this basis, we recommend at least a compromise between the selection of the best model and a model-averaging approach, where point predictions might be derived from the best model and uncertainty be ascertained by BMA.

We chose to highlight one simple approach to incorporating model uncertainty in prediction. There would be

great value in a further detailed study that compares the performance of the various methods available for treating model uncertainty in an ecological context. In particular, a thorough investigation into the value of MCMC, hierarchical, and bootstrap approaches is needed.

The credibility of predictions is of primary importance to conservation decision-makers, who take financial, ecological, and (sometimes) personal risks every time they use the predictions of a model to make a decision. Hedging against uncertainty is an integral part of managing those risks. Conservation biologists who best estimate and communicate the uncertainty of their predictions will be of most use to nervous decision-makers. Model averaging is a logical and simple technique that can provide a more realistic estimate of the uncertainty associated with model predictions.

Acknowledgments

This work was co-sponsored by State Forests of New South Wales, Queensland Parks and Wildlife Service, and the Holdsworth Fund. B.A.W. is particularly grateful to B. Smith, J. Kehl, G. Smith, and B. Holdsworth for their support. We are grateful to S. Bekessy, D. Bardos, M. Burgman, J. Elith, M. Pascual, and one anonymous reviewer for constructive and critical comments.

Literature Cited

- Agresti, A. 1996. An introduction to categorical data analysis. Wiley, New York.
- Akaike, H. 1973. Information theory and an extension of the maximum likelihood principle. Pages 267–281 in B. N. Petrov and F. Csaki, editors. Second international symposium on information theory. Reprinted in S. Kotz and N. L. Johnson, editors. 1992. Breakthroughs in statistics. I. Foundations and basic theory. Springer-Verlag, New York.
- Beissinger, S. R., and M. I. Westphal. 1998. On the use of demographic models of population viability in endangered species management. *Journal of Wildlife Management* **62**:821–841.
- Berger, J. O. 1985. Statistical decision theory and Bayesian analysis. Springer, New York.
- Braithwaite, L. W. 1983. Studies of the arboreal marsupial fauna of eucalypt forests being harvested for woodpulp at Eden, New South Wales. I. The species and distribution of animals. *Australian Wildlife Research* **10**:219–229.
- Braithwaite, L. W., D. L. Binns, and R. D. Nowlan. 1988. The distribution of arboreal marsupials in relation to Eucalypt forest types in the Eden New South Wales Australia wood chip concession area. *Australian Wildlife Research* **15**:363–374.
- Buckland, S. T., and D. A. Elston. 1993. Empirical models for the spatial distribution of wildlife. *Journal of Applied Ecology* **30**:478–495.
- Buckland, S. T., K. P. Burnham, and N. H. Augustin. 1997. Model selection: an integral part of inference. *Biometrics* **53**:603–618.
- Burnham, K. P., and D. R. Anderson. 1998. Model selection and inference: a practical information-theoretic approach. Springer, New York.
- Chatfield, C. 1995. Model uncertainty, data mining and statistical inference. *Journal of the Royal Statistical Society Series A* **158**:419–466.
- Clyde M.A. 1999 Comment on Bayesian model averaging: a tutorial by Hoeting et al. *Statistical Science* **14**:382–417.
- Clyde, M. A. 2000. Model uncertainty and health effect studies for particulate matter. *Environmetrics* **11**:745–763.
- Conroy, M. J., Y. Cohen, F. C. James, Y. G. Matsinos, and B. A. Maurer. 1995. Parameter estimation, reliability, and model improvement for spatially explicit models of animal populations. *Ecological Applications* **5**:17–19.
- Cork, S. J., and P. C. Catling. 1996. Modelling distributions of arboreal and ground-dwelling mammals in relation to climate, nutrients, plant chemical defences and vegetation structure in the eucalypt forests of southeastern Australia. *Forest Ecology & Management* **85**:163–175.
- Derksen, S., and H. J. Keselman. 1992. Backward, forward and stepwise automated subset selection algorithms: frequency of obtaining authentic and noise variables. *British Journal of Mathematical and Statistical Psychology* **45**:265–282.
- Draper, D. 1995. Assessment and propagation of model uncertainty. *Journal of the Royal Statistical Society Series B* **57**:45–97.
- Elith, J., M. A. Burgman, and H. M. Regan. 2002. Mapping epistemic uncertainties and vague concepts in predictions of species distribution. *Ecological Modeling* **157**:313–329.
- Ellison, A. M. 1996. An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecological Applications* **6**:1036–1046.
- Fielding, A. H., and J. F. Bell. 1997. A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environmental Conservation* **24**:38–49.
- Fleishman, E., R. Macnally, J. P. Fay, and D. D. Murphy. 2001. Modeling and predicting species occurrence using broad-scale environmental variables: an example with butterflies of the Great Basin. *Conservation Biology* **15**:1674–1685.
- Foley, W. J., J. C. Kehl, K. A. Nagy, I. R. Kaplan, and A. C. Borsboom. 1990. Energy and water metabolism in free-living greater gliders *Petauroides-Volans*. *Australian Journal of Zoology* **38**:1–10.
- Furnival, G. M., and R. W. Wilson. 1974. Regression by leaps and bounds. *Technometrics* **16**:499.
- George, E. I., and R. E. McCulloch. 1995. Stochastic search variable selection. Pages 203–214 in W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. Markov chain Monte Carlo in practice. Chapman and Hall, London.
- Henry, S. R. 1984. Social organisation of the greater glider (*Petauroides volans*) in Victoria. Pages 221–228 in A. P. Smith and I. D. Hume, editors. Possums and gliders. Australian Mammal Society, Sydney.
- Hoeting, J. A., D. Madigan, A. E. Raftery, and C. T. Volinsky. 1999. Bayesian model averaging: a tutorial. *Statistical Science* **14**:382–401.
- Hurvich, C. M., and C. L. Tsai. 1990. The impact of model selection on inference in linear regression. *The American Statistician* **44**:214–217.
- Hutchinson, M. E., H. A. Nix, D. J. Houlder, and J. P. McMahon. 1999. ANUCLIM: a software package for systematic interrogation of climate surface coefficient files. Version 1.8. Centre for Resource and Environmental Studies, Australian National University, Canberra.
- Jeffreys, H. 1961. Theory of probability. 3rd edition. Oxford University Press, Oxford, United Kingdom.
- Kass, R. E., and A. E. Raftery. 1995. Bayes factors. *Journal of the American Statistical Association* **90**:773–795.
- Kavanagh, R. P. 1997. Ecology and management of large forest owls in south-eastern Australia. Ph.D. dissertation. University of Sydney, Sydney.
- Kavanagh, R. P. 2000. Effects of variable-intensity logging and the influence of habitat variables on the distribution of the greater glider *Petauroides volans* in montane forest, southeastern New South Wales. *Pacific Conservation Biology* **6**:18–30.
- Kavanagh, R. P., and K. L. Bamkin. 1995. Distribution of nocturnal forest birds and mammals in relation to the logging mosaic in south-eastern New South Wales, Australia. *Biological Conservation* **71**:41–53.

- Kavanagh, R. P., and M. J. Lambert. 1990. Food selection by the greater glider, *Petauroides volans*: is foliar nitrogen a determinant of habitat quality? *Australian Wildlife Research* **17**:285–299.
- Keith, D. A., and M. Bedward. 1999. Native vegetation of the South East Forests region, Eden, New South Wales. *Cunninghamia* **6**:1–217.
- Land Information Centre (LIC). 1997. Digital elevation model. LIC, New South Wales State Government, Sydney.
- Li, W., Z. Wang, Z. Ma, and H. Tang. 1999. Designing the core zone in a biosphere reserve based on suitable habitats: Yancheng Biosphere Reserve and the red crowned crane (*Grus japonensis*). *Biological Conservation* **90**:167–173.
- Lindenmayer, D. B., R. B. Cunningham, M. T. Tanton, A. P. Smith, and H. A. Nix. 1990a. Habitat requirements of the mountain brushtail possum and the greater glider in the montane ash-type eucalypt forests of the Central Highlands of Victoria Australia. *Australian Wildlife Research* **17**:467–478.
- Lindenmayer, D. B., R. B. Cunningham, M. T. Tanton, A. P. Smith, and H. A. Nix. 1990b. The conservation of arboreal mammals in the montane ash forests of the Central Highlands of Victoria, south east Australia: I. Factors effecting the occupancy of trees with hollows. *Biological Conservation* **54**:111–131.
- Loyn, R. H., E. G. McNabb, L. Volodina, and R. Willig. 2001. Modelling landscape distributions of large forest owls as applied to managing forests in north-east Victoria, Australia. *Biological Conservation* **97**:361–376.
- Magder, L. S., and J. P. Hughes. 1997. Logistic regression when the outcome is measured with uncertainty. *American Journal of Epidemiology* **146**:195–203.
- MathSoft. 1997. SPLUS 4.5. Data Analysis Products Division, MathSoft, Seattle.
- McCullagh, P., and J. A. Nelder. 1989. *Generalised linear models*. 2nd edition. Chapman and Hall, London.
- McKay, G. M. 1995. *Mammals of Australia*. 2nd edition. Reed Books, Chatswood, Australia.
- Miller, M. E., and S. L. Hui. 1991. Validation techniques for logistic regression models. *Statistics in Medicine* **10**:1213–1226.
- Mladenoff, D. J., and T. A. Sickley. 1998. Assessing potential gray wolf restoration in the northeastern United States: a spatial prediction of favorable habitat and potential population levels. *Journal of Wildlife Management* **62**:1–10.
- Nicholls, A. O. 1989. How to make biological surveys go further with generalised linear models. *Biological Conservation* **50**:51–75.
- New South Wales National Parks and Wildlife Service (NSW NPWS). 1998. *Eden fauna modelling*. New South Wales Regional Forest Agreement Steering Committee, Canberra.
- Pascual, M. A., P. Kareiva, and R. Hilborn. 1997. The influence of model structure on conclusions about the viability and harvesting of Serengeti wildebeest. *Conservation Biology* **11**:966–976.
- Pausas, J. G., L. W. Braithwaite, and M. P. Austin. 1995. Modelling habitat quality for arboreal marsupials in the south coastal forests of New South Wales, Australia. *Forest Ecology & Management* **78**:39–49.
- Pearce, J., and S. Ferrier. 2000. Evaluating the predictive performance of habitat models developed using logistic regression. *Ecological Modelling* **133**:225–245.
- Pereira, J. M. C., and R. M. Itami. 1991. GIS-based modelling using logistic multiple regression: a case study of the Mt. Graham red squirrel. *Photogrammetric Engineering and Remote Sensing* **57**:1475–1486.
- Porter, W. P., S. Budaraju, W. E. Stewart, and N. Ramankutty. 2000. Calculating climate effects on birds and mammals: impacts on biodiversity, conservation, population parameters, and global community structure. *American Zoologist* **40**:597–630.
- Possingham, H. P., D. B. Lindenmayer, T. W. Norton, and I. Davies. 1994. Metapopulation viability analysis of the greater glider *Petauroides volans* in a wood production area. *Biological Conservation* **70**:227–236.
- Queensland Government. 1998. *Systematic vertebrate fauna survey project. Stage II: assessment of habitat quality for priority species in the Southeast Queensland Bioregion*. Queensland Regional Forest Agreement Steering Committee, Canberra.
- Raftery, A. E. 1996. Approximate Bayes factors and accounting for model uncertainty in generalised linear models. *Biometrika* **83**:251–266.
- Raftery, A. E., D. Madigan, and J. A. Hoeting. 1997. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association* **92**:179–191.
- Reichert, P., and M. Omlin. 1997. On the usefulness of overparameterized ecological models. *Ecological Modelling* **95**:289–299.
- Schwarz, G. 1978. Estimating the dimension of a model. *The Annals of Statistics* **6**:461–464.
- Tierney, L., and J. B. Kadane. 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* **81**:82–86.
- Volinsky, C. T., and A. E. Raftery. 2000. Bayesian information criterion for censored survival models. *Biometrics* **56**:256–262.
- Wade, P. R. 2000. Bayesian methods in conservation biology. *Conservation Biology* **14**:1308–1316.

