

Modelling species distributions in Britain: a hierarchical integration of climate and land-cover data

Richard G. Pearson, Terence P. Dawson and Canran Liu

Pearson, R. G., Dawson, T. P. and Liu, C. 2004. Modelling species distributions in Britain: a hierarchical integration of climate and landcover data. – *Ecography* 27: 285–298.

A modelling framework for studying the combined effects of climate and land-cover changes on the distribution of species is presented. The model integrates land-cover data into a correlative bioclimatic model in a scale-dependent hierarchical manner, whereby Artificial Neural Networks are used to characterise species' climatic requirements at the European scale and land-cover requirements at the British scale. The model has been tested against an alternative non-hierarchical approach and has been applied to four plant species in Britain: *Rhynchospora alba*, *Erica tetralix*, *Salix herbacea* and *Geranium sylvaticum*. Predictive performance has been evaluated using Cohen's Kappa statistic and the area under the Receiver Operating Characteristic curve, and a novel approach to identifying thresholds of occurrence which utilises three levels of confidence has been applied. Results demonstrate reasonable to good predictive performance for each species, with the main patterns of distribution simulated at both 10 km and 1 km resolutions. The incorporation of land-cover data was found to significantly improve purely climate-driven predictions for *R. alba* and *E. tetralix*, enabling regions with suitable climate but unsuitable land-cover to be identified. The study thus provides an insight into the roles of climate and land-cover as determinants of species' distributions and it is demonstrated that the modelling approach presented can provide a useful framework for making predictions of distributions under scenarios of changing climate and land-cover type. The paper confirms the potential utility of multi-scale approaches for understanding environmental limitations to species' distributions, and demonstrates that the search for environmental correlates with species' distributions must be addressed at an appropriate spatial scale. Our study contributes to the mounting evidence that hierarchical schemes are characteristic of ecological systems.

R. G. Pearson (richard.pearson@eci.ox.ac.uk), T. P. Dawson and C. Liu, *Environmental Change Inst., School of Geography and the Environment, Univ. of Oxford, 1a Mansfield Road, Oxford, U.K. OX1 3SZ.*

Climate change and land-cover change, leading to habitat loss, are two of the most important factors threatening ecosystems worldwide (Sala et al. 2000, Parmesan and Yohe 2003). Considered individually, the threats posed are significant, but the interaction between the two factors could be disastrous (Travis 2003). A number of studies have modelled the potential impacts of climate change on the distribution of species, applying climate-driven simulations across a range of scales, study

areas and species (for reviews see Guisan and Zimmermann 2000, Pearson and Dawson 2003). However, few modelling studies have explicitly addressed climate – land-cover interactions. In order to study the combined effects of climate change and habitat fragmentation (as driven by changes in land-cover), it is necessary to develop a modelling framework that integrates climate and land-cover drivers. This paper presents a novel modelling approach that incorporates fine-scale land-

Accepted 12 December 2003

Copyright © ECOGRAPHY 2004
ISSN 0906-7590

cover data into a static, correlative model driven by coarse-scale climate data. The model has been applied to four plant species in Britain, providing an insight into the roles of climate and land-cover as correlates of these species distributions.

The spatial scale at which species distribution modelling is undertaken is of fundamental importance, with the selection of appropriate spatial extent and data resolution (or “grain”) for a given application essential. It has been proposed that climate impacts on the distribution of species will be most apparent at macro-scales, with broad spatial extents and coarse data resolutions most appropriate for correlating climate with species distributions (Pearson et al. 2002). It has also been hypothesised that within the climate space defined by synoptic climate conditions other factors influence the distribution of species in a hierarchical manner, with different factors being better correlates at different scales (Franklin 1995, Collingham et al. 2000, Pearson and Dawson 2003). Thus, it is proposed that land-cover may be considered the dominant control of species presence and absence at a finer spatial resolution than climate. This concept has been applied and tested in this study through the development of a modelling methodology that integrates climate and land-cover data in a scale-dependent hierarchical manner.

The incorporation of fine-scale land-cover data into a model designed to simulate broad-scale climate suitability is an important methodological development. The refined model aims to identify areas suitable for a given species at a finer resolution than in purely climate-driven studies. Regional-scale predictions of potential future environmental impacts will be more suited to the requirements of local conservation policy planning than those of previous studies looking at the potential impacts and policy implications of broad-scale climate change (e.g. Berry et al. 2002, 2003). The downscaling of climate-driven modelling will also facilitate coupling of the modelled species’ suitability surfaces with dynamic simulations of species dispersal (e.g. Collingham et al. 1996). Such simulations are commonly parameterised to operate over much finer-resolution landscapes than those generated by the climate-driven models.

In this context, this study aimed to design and test a modelling methodology for generating regional fine-scale suitability surfaces for species. A suitability surface is defined as a landscape identifying areas where a species could potentially grow and reproduce, and is analogous to an approximation of the spatial manifestation of the fundamental niche. The model has been designed to be applicable to scenarios of future climate and land-cover change. A hierarchical framework of environmental controls on the distribution of species has been applied and the utility and validity of this approach has been tested.

Methods

General Modelling Principles

The generation of fine-scale suitability surfaces has been achieved through modification of the bioclimatic model “SPECIES” (Pearson et al. 2002). SPECIES employs an Artificial Neural Network (ANN) to characterise bioclimate envelopes based on inputs generated from climate and soils data using a climate-hydrological process model. An important element of the model is its multi-scale approach whereby the bioclimate envelope of a species is first identified at a coarse resolution (0.5°) European extent before application at a finer-resolution British extent. The aim is to characterise the species-climate relationship at the macro-scale, where climate impacts are expected to be most apparent (Pearson and Dawson 2003). This multi-scale approach enables climatic range margins that are currently outside Britain, but which may move into Britain under future climate scenarios, to be identified. The model therefore does not extrapolate outside its training dataset when used to predict the distribution of bioclimate envelopes in Britain under potential future climates.

The SPECIES model has been extended to incorporate land-cover inputs through a modelling process as outlined in Fig. 1. The top half of the schematic presents the original SPECIES model, with continental-scale ANN training driven by climate. Land-cover data at 10 km resolution is then incorporated along with the British-scale bioclimate suitability surface as input into a second ANN that is trained against British species distributions. This second ANN, which aims to characterise the relationship between species distribution, climate and land-cover, has then been used to simulate potential distributions at 10 km and 1 km resolutions. The climate suitability surface is thus refined based on correlations between land-cover type and observed distribution. In this way a hierarchical modelling process is established whereby environmental inputs are incorporated at different spatial scales.

In order to test the hierarchical model an alternative approach has also been applied in this study. Hereby, the European-scale modelling was discarded and, instead, the climate and land-cover inputs were both correlated against the observed distribution at only the 10 km resolution and British extent. Theoretically, these non-hierarchical simulations should not be so robust under scenarios of future climate change since the broader bioclimate envelope is not considered. However, comparison between the abilities of the two modelling approaches to simulate current distributions has enabled the validity of the hierarchical scheme to be examined.

Four plant species, representing a range of life forms, habitat associations and distribution characteristics, were selected for modelling. Suitable climate space for each of the species had been modelled previously using

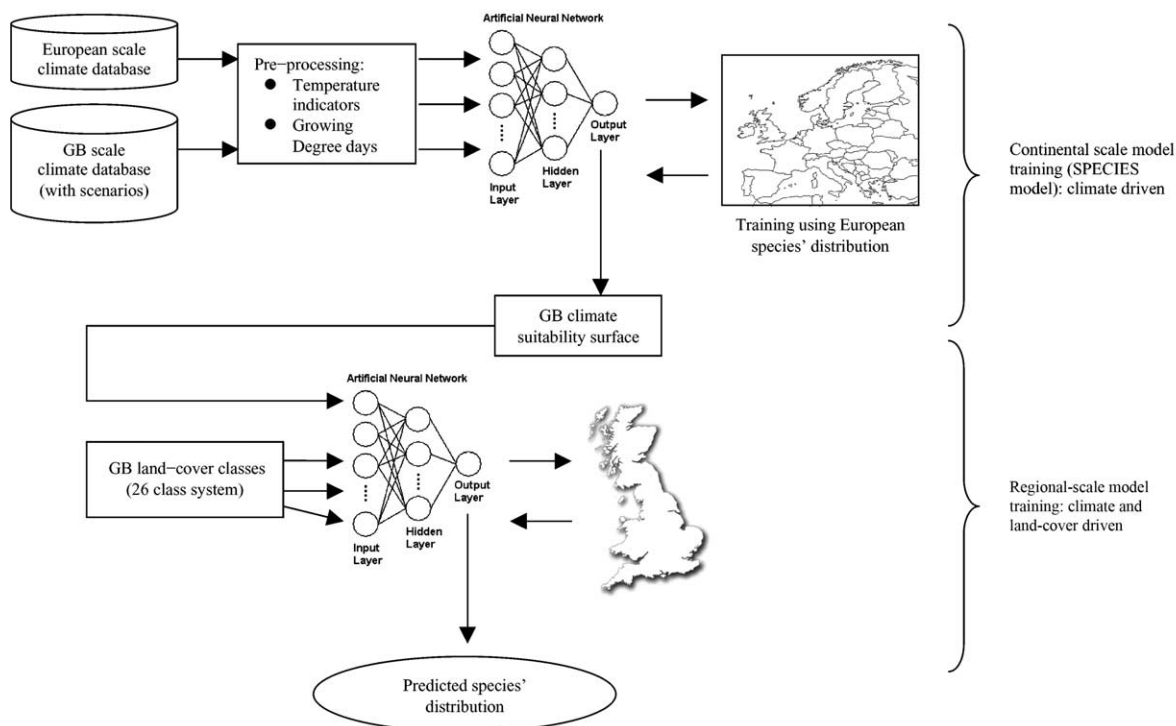


Fig. 1. Schematic of the hierarchical modelling process.

the SPECIES model (Pearson et al. 2002). Observed distributions and simulated climate suitabilities are presented in the first two columns of Fig. 2. *Rhynchospora alba* (white beak-sedge) is a perennial herb characteristic of lowland acidic bogs, wet heaths and mires. The species has a disjunct distribution throughout Britain and has a modelled climate envelope that covers much of southern England and the Scottish highlands. *Erica tetralix* (cross-leaved heath) is a low shrub found in a wide range of mires and wet heaths, extending into drier heath in south west Britain. The species is common throughout much of Britain but its distribution is fragmented in southern England. SPECIES simulations suggest that *E. tetralix* has suitable climate space throughout Britain. *Salix herbacea* (dwarf willow) is a dwarf shrub found under conditions of extreme exposure and characteristic of Arctic-montane habitats. *Salix herbacea* is found in the Scottish Highlands, Cumbria and in upland parts of Wales. The simulated climate space for this species identifies these upland regions but has the tendency to extend further south than the observed distribution. *Geranium sylvaticum* (wood crane's-bill) is a perennial herb of upland hay meadows, ungrazed damp woodlands, streamsides and mountain rock ledges, and is often found on lane-sides. *Geranium sylvaticum* is native to much of northern Britain, though is uncommon in the far north of Scotland. The species has a number of isolated presences south of its main

range due largely to alien introductions (Preston et al. 2002). The simulated climate space for *G. sylvaticum* covers much of northern Britain and extends south down the Pennine ridge and throughout much of Wales.

Datasets and model inputs

Species distributions

The study utilised presence/absence species distributions at three different spatial extents/resolutions: European extent at 0.5° resolution, British extent at 10 km resolution, and county extent at 2 km resolution. European distributions were obtained from several European atlases and used for ANN training as described in Pearson et al. (2002). British distributions were obtained in digital format from the Biological Records Centre (Centre for Ecology and Hydrology, Monks Wood, U.K.). All post-1930 recorded observations, both native and alien, were included in the analysis. Three regions of Britain were chosen for testing the 1 km resolution simulations: Cumbria, Devon and East Anglia. These regions were selected based on the availability of local floras and so as to give a range of species prevalences. Local distributions were digitised from atlases of county floras for Cumbria (Halliday et al. 1997), Devon (Ivimey-Cook 1984) and Norfolk (Beckett et al. 1999). Some additional records were

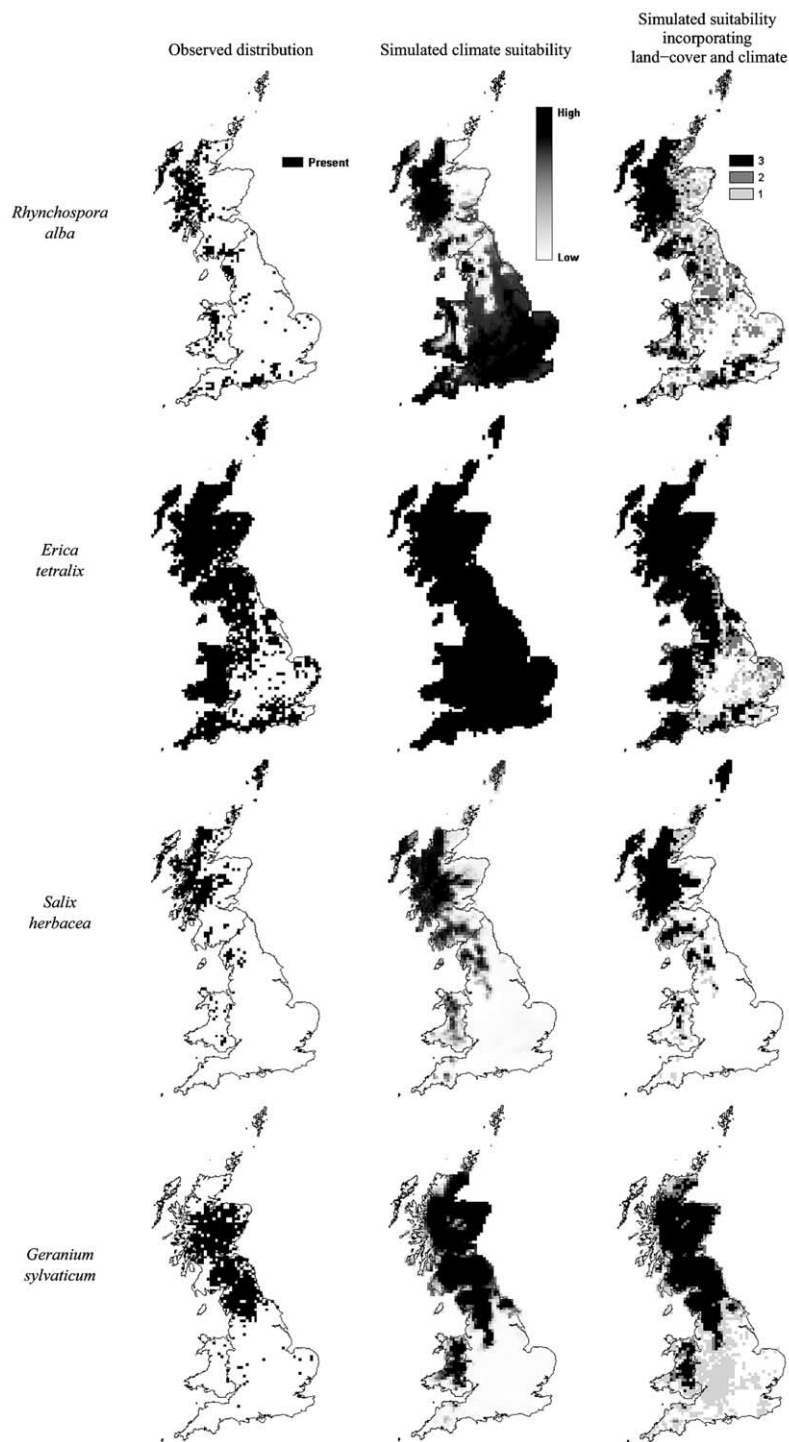


Fig. 2. Observed and simulated distributions in Britain at 10 km resolution for the four species tested. Observed distributions were obtained from the Biological Records Centre. Simulated climate suitabilities were modelled using the original SPECIES model (i.e. climate envelope identified at the continental scale; 0.5° resolution). The right hand column shows how these climate space simulations can be further refined by incorporating land-cover inputs into the model processing at the regional scale (10 km resolution). Simulations incorporating both land-cover and climate are presented at three thresholds as described in the text: 3 = sensitivity/specificity balanced threshold; 2 = 90% sensitivity; 1 = 95% sensitivity.

obtained in digital format from Suffolk Biological Records Centre (Ipswich Museum). Regional distributions were available at 2 km (tetrad) resolution and were plotted as the centre-point of four 1 km cells.

Climate and soils

Climate and soils data were utilised in the study at both the 0.5° European scale and the 10 km British scale. Details regarding the data sources and processing for the

climatic element of this study are described in full in Pearson et al. (2002). Briefly, climatologies for the 1961–1990 normal (Hulme et al. 1995, Hulme and Jenkins 1998) were processed along with an estimation of soil water-holding capacity to generate five bioclimatic inputs which are thought to have direct physiological roles in limiting plant ranges. The bioclimatic inputs were: absolute minimum temperature expected over a 20 yr period, maximum annual temperature, growing degree days $> 5^{\circ}\text{C}$, annual soil moisture surplus and annual soil moisture deficit. These inputs were used to generate a single bioclimate suitability surface for Britain, scaled from 0 to 1, for each species through the multi-scale SPECIES approach. This suitability surface was at a resolution of 5 km due to the availability of soils data at this resolution. In order to match the resolution of the British species distributions it was necessary to artificially aggregate the suitability surface to 10 km resolution by taking the mean of each block of four 5 km cells. This 10 km British climate suitability surface was used as model input alongside land-cover data in the hierarchical modelling approach described in this study. The alternative non-hierarchical methodology tested in the study utilised the five raw bioclimatic variables as ANN inputs at the British scale only. For the non-hierarchical approach the bioclimatic inputs were normalised using a linear transformation based on the minimum and maximum values for Britain so as to ensure that the ranges of the inputs did not differ by orders of magnitude (Tarassenko 1998).

Land-cover

Land-cover data for Britain were provided by the Land Cover Map 2000 (LCM2000) in which the distribution of 26 terrestrial classes of broad habitat types (subclass level 2; for details see www.ceh.ac.uk/data/lcm/lcm2000.shtm) is recorded. Data was obtained from the Centre for Ecology and Hydrology, Monks Wood, U.K., as grid-based percentage cover 1 km^{-2} . The presence of diagonal banding in this dataset that matched the path of landsat imagery required that four land-cover classes be removed from the analysis (setaside grass, arable cereals, arable non-rotational and inland bare ground). The remaining percentages for the 22 land-cover types were reformatted to give presence/absence data for each class at both 10 km and 1 km resolutions. This 10 km dataset was used alongside the British climate suitability surface to correlate against the observed British species distribution, whilst the 1 km dataset was subsequently used to simulate distributions at the finer resolution.

Artificial Neural Networks

ANNs are computer systems that have increasingly been employed in ecological studies as an alternative to more

traditional statistical techniques (Lek and Guegan 1999). Inspired by the structure of the brain, ANNs consist of many processing elements (artificial neurons) that are interconnected to form a network. ANNs are “trained” by repeatedly passing large numbers of known examples of the problem under consideration through the network. By repeatedly adjusting the connections between processing elements the difference between the network predictions and the known examples can be minimised. A successfully trained ANN can then be used to make predictions for inputs where the output is not known (Tarassenko 1998).

The advantages and disadvantages of using ANNs for characterising species distributions have been discussed in detail by Hilbert and Ostendorf (2001) and Pearson et al. (2002). Of particular note here is the ability of ANNs to identify non-linear responses to environmental variables and to incorporate multiple types of input variables, including categorical (e.g. land-cover classes) and non-categorical (e.g. climate suitability) data. A notable disadvantage is that the relative contribution of different input variables is not immediately identified in an ANN, though further analysis of the network can increase the explanatory power of the approach (Gevrey et al. 2003).

This study used feed-forward ANNs with node connection weights adjusted after every training pattern using a backpropagation learning algorithm (generalised delta rule). Network training employed a sigmoidal activation function and a quadratic (mean squared) error function (Maier and Dandy 1998). ANNs were constructed and trained using SNNS (Stuttgart Neural Network Simulator) ver. 4.1 software (Univ. of Stuttgart, Germany: <http://www-ra.informatik.uni-tuebingen.de/SNNS>). Network architecture and parameter values for the British scale modelling were decided upon following extensive tests to optimise performance. For the hierarchical approach networks were constructed with 23 input nodes (for 22 land-cover classes and 1 climate suitability surface), 47 nodes in a single hidden layer, and one output node (for species presence/absence). For the non-hierarchical approach networks with 27 input nodes (22 land-cover classes and 5 bioclimatic variables), 55 nodes in a single hidden layer, and one output node were used. Initial connection weights were selected randomly in the range 1.0 to -1.0 , and the learning parameter set at 0.05. Use of a momentum term, which controls the amount by which weights are adjusted in each iteration so as to avoid local minima, was found to improve model training and was included with a value of 0.5.

For each British scale model run (ie. hierarchical and non-hierarchical models for each of the four species) the full 10 km dataset was randomly split into three parts for network training (1387 points), validating (694 points) and testing (693 points). The validation dataset was used to monitor network error after each training cycle,

enabling training to be stopped when the network began to over-train (Tarassenko 1998). For each model run ten networks, each with different weight initialisations, were trained and the network achieving the lowest minimum error for the validation dataset was selected. The test dataset remained “unseen” by the ANN during network training and was used to evaluate the predictive performance of the model.

Evaluating predictive performance

A number of methods for assessing the predictive performance of presence/absence models in ecology have been proposed (for reviews see Fielding and Bell 1997, Pearce and Ferrier 2000). This study utilised two tests: Cohen’s Kappa statistic of similarity (k) and the Area Under the Receiver Operating Characteristic Curve (AUC).

Kappa is a commonly used statistic that provides a measure of proportional accuracy, adjusted for chance agreement (Cohen 1960, Monserud and Leemans 1992). Though providing a straightforward measure of agreement, k has two notable limitations. Firstly, it is affected by prevalence, making comparison of values between different species and different studies difficult (Fielding and Bell 1997). Secondly, k is dependent on the application of a decision threshold (above which model outputs are considered to represent “presence”) and therefore does not measure predictive accuracy across the full range of possible thresholds. However, by calculating k across a range of thresholds, its value can be maximised and an “optimum” threshold, representing maximum agreement between observed and simulated distributions, can be identified. The maximised k has been calculated in this study.

In view of the limitations of k and other discrimination indices, Fielding and Bell (1997) concluded that AUC is the index that best meets the requirements of an unbiased measure of accuracy. AUC is calculated from the Receiver Operating Characteristic (ROC) curve. The ROC curve describes the compromise that is made between the sensitivity (defined as the proportion of true positive predictions vs the number of actual positive sites) and false positive fraction (the proportion of false positive predictions vs the number of actual negative sites) as the decision threshold is varied. The curve is defined by varying the decision threshold incrementally across the range of predicted model outputs and generating a series of pairs of sensitivity and false positive fraction values that can be plotted on a graph (Pearce and Ferrier 2000). Swets (1988) proposed that the best discrimination index is the area under the ROC curve expressed as a proportion of the total area of the unit square defined by the axes. This index is independent of both species prevalence and the decision threshold. AUC ranges from 0.5 for models with no

discrimination ability, to 1 for models with perfect discrimination.

AUC was calculated using the method described by Hanley and McNeil (1982) which is based on the derivation of the Wilcoxon statistic. Pearce and Ferrier (2000) recommend this method for ecological applications since it is nonparametric. The standard error of AUC was estimated by bootstrapping with a sample of 1000, as recommended by Pearce and Ferrier (2000).

Despite the theoretical advantages of AUC, Manel et al. (2001) found that in practice k was highly significantly correlated with AUC and concluded that k provides a simple and effective test of predictive performance. Since k has been most widely applied to date, and in order to investigate further the relationship between k and AUC, both measures of predictive performance have been used in this study. The statistics were calculated for model predictions made on the “unseen” test dataset.

Identifying decision thresholds

To aid the interpretation and presentation of model results it is useful to identify a threshold value above which model outputs are considered to represent species presence (or to define that the site is suitable for habitation). The choice of threshold value is important because model outputs, when mapped as presence/absence, may look quite different dependent on the threshold applied. A threshold is commonly identified by maximising the agreement between observed and simulated distributions. One approach is to use the threshold value that maximises the Kappa statistic of agreement, as described in the previous section (and applied in Pearson et al. 2002). An alternative approach, based on the ROC procedure, is to plot sensitivity against specificity (defined as the proportion of true negative predictions vs the number of actual negative sites) at a series of thresholds and to apply the threshold value at which these two curves cross (Fig. 3, threshold a; Thuiller et al. 2003). In this way, the cost arising from an incorrect decision is balanced against the benefit gained from a correct prediction (Manel et al. 2001).

The above threshold values, based on maximising agreement between observed and simulated distributions, have been calculated in this study. However, it may be argued that these approaches do not represent the most appropriate threshold. Firstly, the choice of a threshold is influenced by the application of the model. For example, if the model is intended to identify potential introduction sites for an endangered species, then it is important to reduce the risk of failure and choose a relatively high threshold that would result in identification only of those sites with a high suitability. In contrast, if the aim is to identify areas within which disturbance may impact the species (e.g. as part of an




Fig. 3. Sensitivity and specificity plotted against threshold for defining decision thresholds. Threshold *a* is assigned at the point where the two curves cross. Thresholds *b* and *c* are defined by sensitivities of 90 and 95% respectively. The curves presented are those for the simulation of *Erica tetralix*.

environmental impact assessment or for the designation of a nature reserve) then the threshold would be set low so as to identify all potentially suitable habitats (Pearce and Ferrier 2000). A second argument against simply maximising agreement concerns the basic aim of the modelling approach presented: to identify those sites where a species could exist (i.e. the potential distribution). Given the many factors that influence the actual distribution of species, maximising the fit between simulated suitabilities and the observed distribution is likely to result in an underestimation of the extent of the potential distribution. In fact, we should be more concerned with minimising the number of sites with observed presences that are predicted to be unsuitable, than with minimising the number of sites without actual presences that are simulated as suitable. More formally, the primary concern when identifying a decision threshold should be to minimise the false negative fraction (the proportion of false negative predictions vs the number of actual positive sites).

Minimising the false negative fraction is analogous to maximising sensitivity (since $\text{sensitivity} = 1 - \text{false negative fraction}$). It is therefore appropriate to define thresholds by assigning cut-off sensitivities as outlined in Fig. 3. Cut-off values defined by sensitivity values of 90 and 95% have been applied in this study (thresholds *b* and *c* in Fig. 3). These thresholds are presented along with those delimited by balancing sensitivity and specificity, giving three levels of confidence in model simulations. Thresholds were defined based on model simulations for the entire British dataset.

Results

10 km resolution simulations

Figure 2 presents simulated distributions generated using the hierarchical model (incorporating land-cover and

climate) alongside the observed distributions and the modelled climate suitabilities. AUC and maximised *k* statistics comparing agreement between the observed distributions and hierarchical simulations are presented in Table 1 (first two columns). It is apparent that the hierarchical simulations are generally good, with the main patterns of distribution identified in each case. Both AUC and Kappa values can be interpreted as indicating reasonable to good model discrimination ability according to the subjective guidelines of Swets (1988) and Monserud and Leemans (1992). Thresholds of occurrence identified for each species are given in Table 2. Differences between the sensitivity/specificity defined thresholds and those identified by maximising Kappa are of note, emphasising the difficulties associated with identifying decision thresholds.

Visual comparison between simulated climate suitabilities and simulated distributions incorporating land-cover data in Fig. 2 demonstrates significant improvement in model performance for the species *R. alba* and *E. tetralix*. For *R. alba* the original bioclimatic model simulated suitable climate space throughout much of central and southern England – regions where the species is in fact scarce – yet the refined model identified the distribution much more accurately by removing those locations where suitable land-cover is not found. Thus, much of central England is identified as having unsuitable combinations of climate and land-cover type, whilst main population centres such as those on the south coast are recognised. Similarly, for *E. tetralix* the SPECIES model identified suitable climate throughout Britain, whilst the downscaled model incorporating land-cover correctly simulated the fragmentation of the distribution in much of central England. Results for *S. herbacea* and *G. sylvaticum* do not show the same degree of improvement between the SPECIES model results and the downscaled methodology.

Turning to the comparison of hierarchical and non-hierarchical approaches, Table 1 presents AUC and maximised *k* values for models generated using each approach. It is apparent that neither method performs consistently better than the other. For *S. herbacea* the hierarchical methodology gave slightly better results according to both AUC and *k* statistics, whilst for *E. tetralix* and *G. sylvaticum* the non-hierarchical approach was superior. In the case of *R. alba*, AUC and *k* gave conflicting results as to which model was more accurate. Differences in both AUC and *k* between the modelling approaches was greatest for *E. tetralix*. Mapping the results for *E. tetralix* (Fig. 4) shows that simulations using the hierarchical and non-hierarchical approaches were very similar, indicating that the statistical differences between simulations can be considered small. This finding was consistent for the other three species and it was therefore concluded that the hierarchical methodology, which has theoretical advantages for the simulation

Table 1. Testing model performance: AUC and maximised Kappa values for independent test datasets. Results are presented for models trained using continental scale climate data with regional land-cover data (hierarchical model), and for models trained with regional scale climate and land-cover data (non-hierarchical model).

	Hierarchical model		Non-hierarchical model	
	AUC (SE)	Maximised Kappa	AUC (SE)	Maximised Kappa
<i>Rhynchospora alba</i>	0.864 (0.019)	0.480	0.845 (0.021)	0.528
<i>Erica tetralix</i>	0.867 (0.014)	0.593	0.900 (0.012)	0.655
<i>Salix herbacea</i>	0.938 (0.012)	0.639	0.916 (0.016)	0.581
<i>Geranium sylvaticum</i>	0.881 (0.017)	0.579	0.913 (0.011)	0.629

of future distributions, does not reduce the quality of simulations of current British distributions.

1 km resolution simulations


ANNs trained against the 10 km resolution British species distribution using the hierarchical modelling approach were used to simulate 1 km resolution suitability surfaces for Britain. Data inputs were 1 km resolution presence/absence land-cover data and 10 km resolution climate data, as described above. Figure 5 presents the simulated suitability surface for *E. tetralix*, with the simulation tested against observed local distributions in the three sub-regions (Cumbria, Devon and East Anglia). At the British extent, the 1 km simulated suitability surface identified the main patterns of distribution evident in the 10 km observed distribution (Fig. 2), but with fragmentation within these limits as expected at the finer resolution. Looking more closely at the sub-regions, it is evident that the model captured the broad distributional trends, though accuracy was reduced at this resolution.

Figure 6 presents 1 km resolution simulations with observed Cumbrian distributions for *R. alba*, *S. herbacea* and *G. sylvaticum*. At the British extent patterns of distribution for each species were similar, though more fragmented, to those in the 10 km observed dataset (Fig. 2). For *G. sylvaticum* it is noted that the simulation retains some of the coarser patterns evident in the 10 km climate data, suggesting that the model for this species primarily reflected climatic limits. Focusing on Cumbria, it is apparent that for each species the simulated suitability surface captured the main patterns of distribution, supporting the proposal that the model is able to successfully estimate potential areas of species occurrence.

Discussion

The model presented shows good potential for integrating the effects of land-cover into an existing bioclimate modelling framework. The interaction between climate and habitat availability plays an important role in determining the biogeography of species and the hierarchical modelling framework presented goes some way to deepening our understanding of this interaction.

Taking the example of *E. tetralix*, Preston et al. (2002) note that this species was mapped as present throughout Britain in the 1962 Atlas of the British Flora (Perring



Hierarchical model simulation Non-hierarchical model simulation

Fig. 4. Comparison of model simulations for *E. tetralix* using a model trained using continental scale climate data with regional land-cover data (hierarchical model), and with regional scale climate and land-cover data (non-hierarchical model). Simulations are presented at three thresholds as described in the text: 3 = sensitivity/specificity balanced threshold; 2 = 90% sensitivity; 1 = 95% sensitivity.

Table 2. Thresholds of occurrence for each of the species modelled identified by four methods, as described in the main text.

	Sensitivity/specificity balance	Sensitivity at 90%	Sensitivity at 95%	Kappa maximisation
<i>Rhynchospora alba</i>	0.11	0.05	0.02	0.30
<i>Erica tetralix</i>	0.73	0.56	0.34	0.63
<i>Salix herbacea</i>	0.16	0.13	0.04	0.29
<i>Geranium sylvaticum</i>	0.35	0.15	0.03	0.51

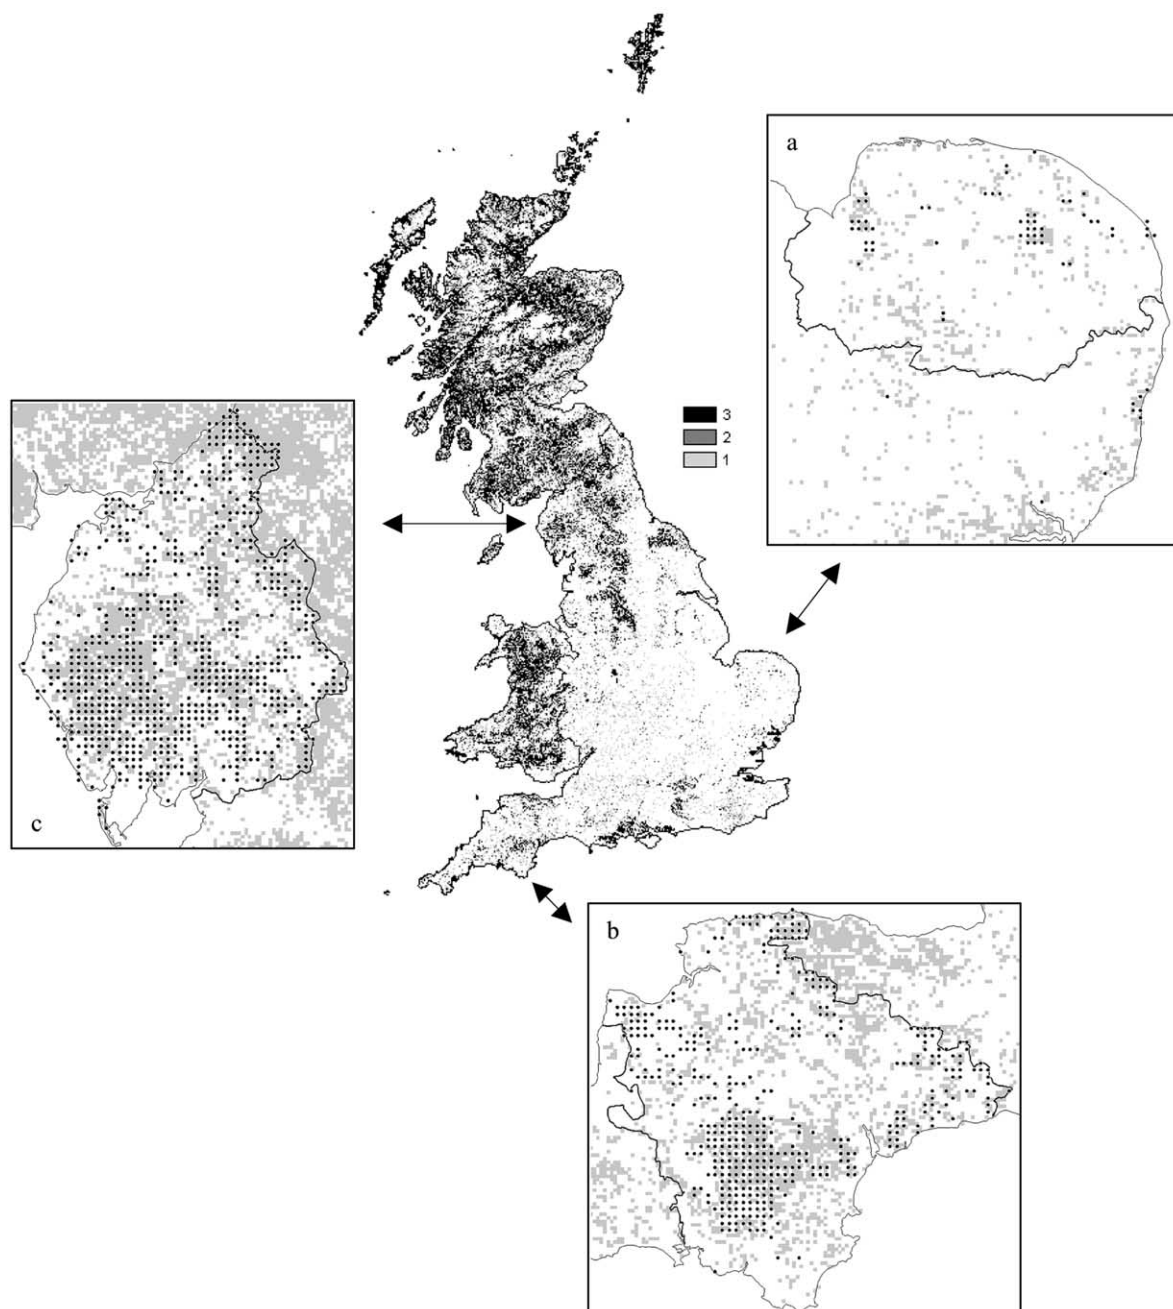


Fig. 5. Simulated U.K. suitability surface for *E. tetralix* at 1 km resolution, with windowed observed presences for county floras (a, Norfolk and Suffolk; b, Devon; c, Cumbria). The full U.K. suitability surface is presented at three thresholds, identified at 10 km resolution, as described in the main text: 3 = sensitivity/specificity balanced threshold; 2 = 90% sensitivity; 1 = 95% sensitivity. Suitability surfaces for the county windows are presented at the 95% sensitivity threshold (grey) along with the observed tetrad distributions (black spots) from county floras (Halliday et al. 1997, Ivimey-Cook 1984, Beckett et al. 1999, M. Sanford pers. comm.).

and Walters 1962) and that the widespread decline in central England since then is most likely attributable to the loss of suitable habitat. Model results presented here support this conclusion and offer no support for the alternative possibility that climate warming since the

mid-twentieth century has caused the decline. Simulations suggest that suitable climate space remained present throughout Britain in the second half of the century (climate data was for the 1961–1990 mean) and that it is the availability of suitable land-cover that is

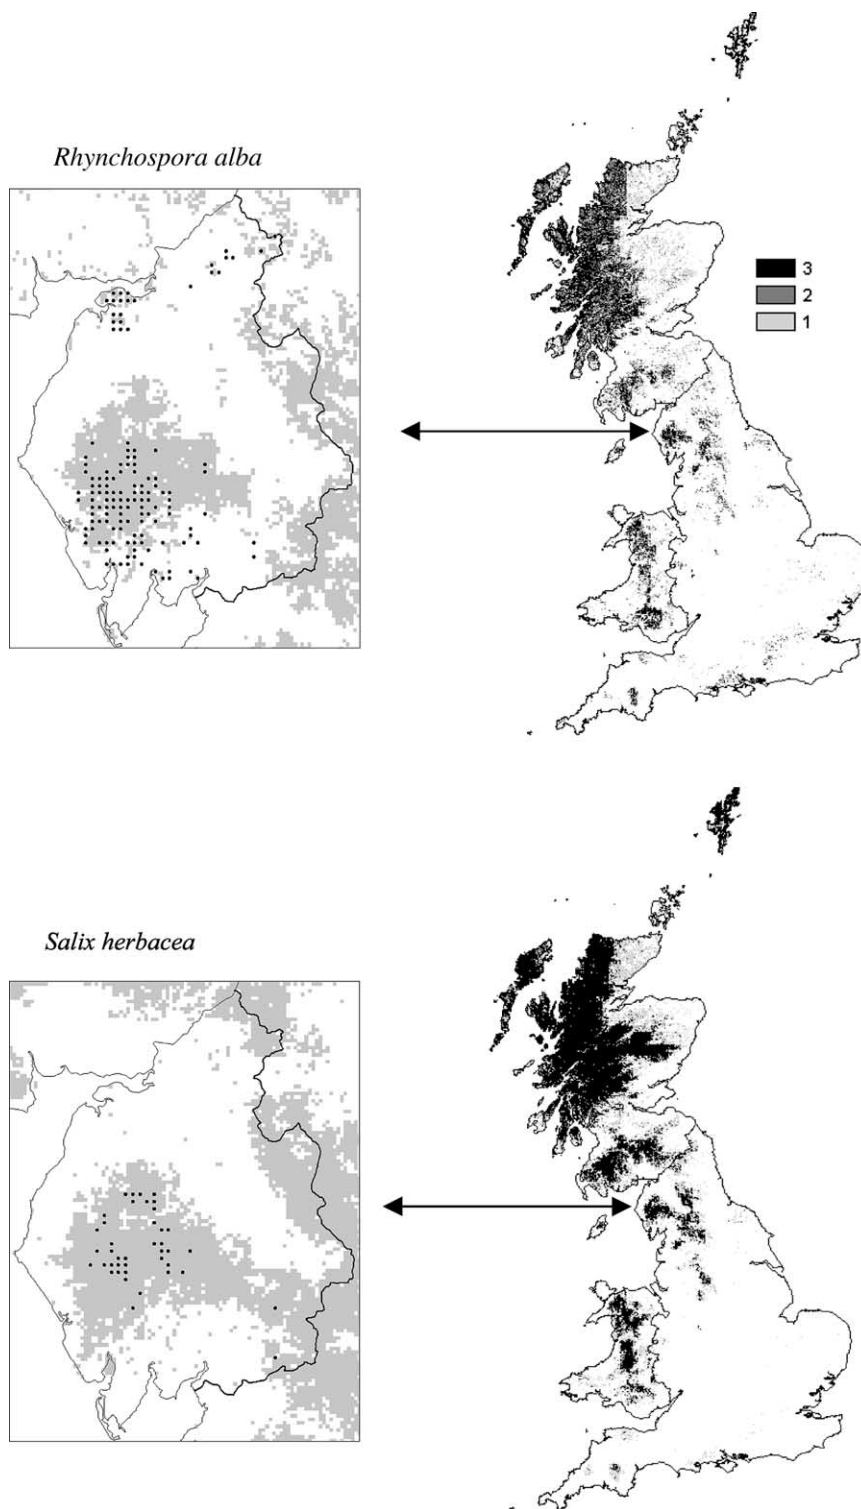
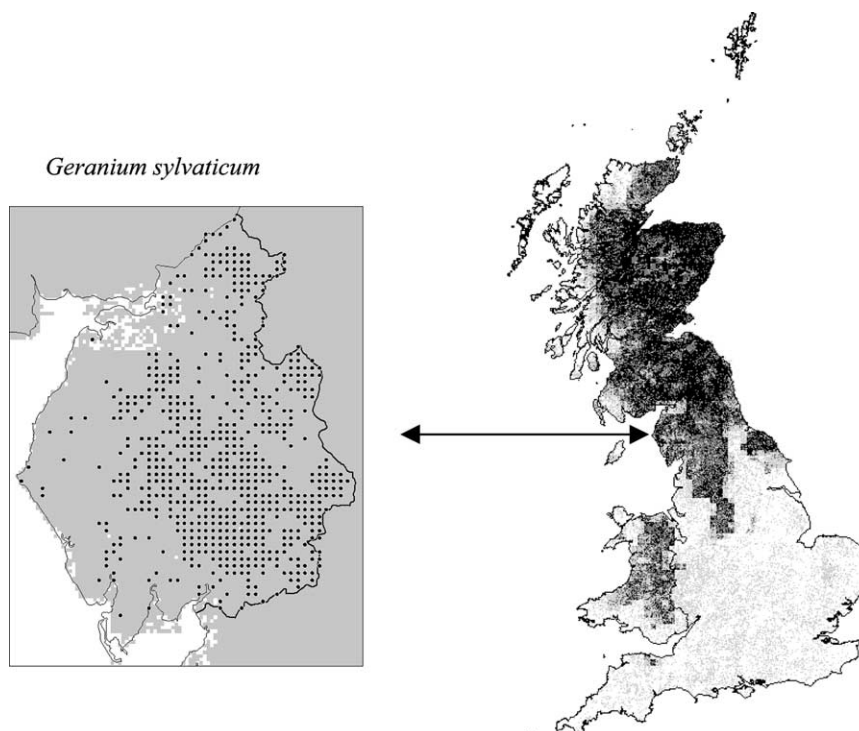


Fig. 6. Simulated U.K. suitability surfaces at 1 km resolution, with windowed observed presences for Cumbria. The full U.K. suitability surfaces are presented at three thresholds, identified at 10 km resolution, as described in the main text: 3 = sensitivity/specificity balanced threshold; 2 = 90% sensitivity; 1 = 95% sensitivity. Suitability surfaces for the county windows are presented at the 95% sensitivity threshold (grey) along with the observed tetrad distributions (black spots) from the Flora of Cumbria (Halliday et al. 1997).

restricting the current species' range. The modelling approach thus has the potential to help uncouple effects of climate and habitat change in the interpretation and

prediction of distribution changes. The importance of discriminating between forces of climate and land-cover change has been demonstrated by Warren et al. (2001) in

Fig. 6 (Continued).



their study of distribution changes of British butterflies over the past 30 yr. It was shown that mobile, habitat generalist species with northern range margins in Britain have increased their ranges in response to climate warming, whilst habitat specialists have declined as a result of habitat loss.

Models for *S. herbacea* and *G. sylvaticum* did not show the same predictive improvements as the other two species when land-cover data was incorporated. It is apparent, therefore, that the presence/absence of land-cover types at 10 km resolution does not provide a good correlate with distribution for these two species. This is likely to be due to the fact that at this resolution nearly all 10 km² cells incorporate at least one small patch of suitable land-cover (i.e. a “presence”), leading to blanket coverage throughout the study region. In order to better identify correlations between land-cover type and distribution of these species it would be necessary to adopt a finer resolution of analysis at which patterns in the distribution of suitable land-cover are apparent in the dataset. These results support the proposal that the search for environmental correlates with species distributions must be addressed at an appropriate spatial scale, whereby distribution patterns for environmental variables and species distributions are similar at the resolution of analysis.

Differences between decision thresholds identified by maximising *k* and by balancing sensitivity and specificity were found to be significant, ranging from 0.10 to 0.19

for the four species studied. The arbitrariness of this threshold must be considered when interpreting presence/absence maps generated from such models. It has been argued here that, rather than maximising agreement between observed and simulated distributions, a more appropriate approach to identifying decision thresholds is to minimise the number of observed presences falling outside the simulated distribution. The three-level approach to presenting model output applied in this study has made the interpretation of results less dependent on the choice of a single threshold and has facilitated the identification of broader potential distributions. For example, the simulated suitability surface for *G. sylvaticum* identified areas south of the main range margin where the potential distribution is simulated at the 90% sensitivity level. The small number of observed presences within these areas (mostly alien introductions) confirms the validity of identifying these regions as having climate and land-cover characteristics that fall within the species’ fundamental niche.

The identification of potential distributions has been a goal of this study. A modelling methodology has been presented based on the hypothesis that hierarchical levels can be identified whereby different environmental drivers have dominant roles in defining species distributions at different spatial resolutions. Relationships between distributions and environmental variables have been characterised at one scale, assuming a degree of equilibrium, and then the relationships identified have been down-

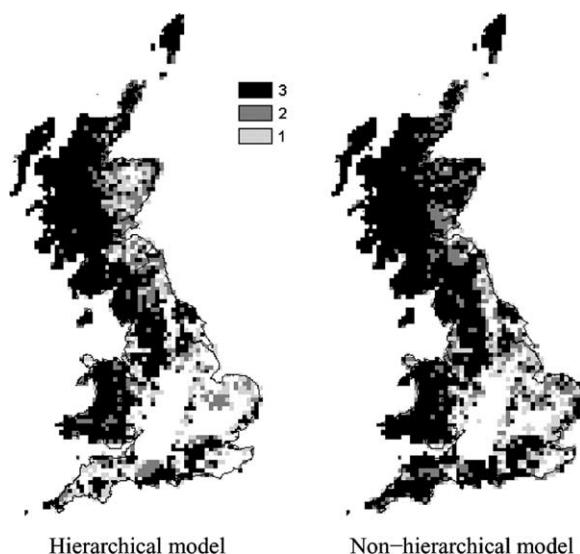


Fig. 7. Comparison of model simulations showing changes to the potential distribution of *Rhynchospora alba* under a high emissions climate change scenario for the 2050s. The left panel presents results using a model trained with continental scale climate data and regional land-cover data (hierarchical model). The right panel presents results using a model trained with regional scale climate and land-cover data (non-hierarchical model). Simulations are presented at three thresholds as described in the text: 3 = sensitivity/specificity balanced threshold; 2 = 90% sensitivity; 1 = 95% sensitivity.

scaled and combined with other environmental factors at a finer spatial scale. The application of this approach to climate data at the European extent and 0.5° resolution, and to land-cover data at the British extent and 10 km resolution, has yielded promising results and it is proposed that the approach defines a useful framework for addressing the complexity of the natural system.

Conclusive validation of the hierarchical scheme is problematic, being confounded by data limitations (notably the lack of species distribution data at a large extent and fine resolution) and by fundamental issues of scale (not least the Modifiable Aerial Unit Problem; Openshaw 1984). Collingham et al. (2000) found no evidence of a hierarchy of environmental controls in their study of invasive weeds across resolutions of 2×2 and 10×10 km. It is possible that similar analysis across a broader range of scales, incorporating 0.5° resolution data, would have generated different results. However, there is mounting evidence that scale-dependent hierarchies of environmental controls are apparent in a number of ecological systems (e.g. Rettie and Messier 2000, Martinez et al. 2003). This study has confirmed the potential utility of the hierarchical approach and it is reaffirmed that the search for different levels of organisation in ecological systems is of potential importance (May 1989).

The modelling approach presented will be particularly useful for estimating potential future impacts of climate change on species' distributions. By formulating the bioclimatic model at the European extent, the model is not required to extrapolate outside its training dataset when predicting future distributions in Britain under climate change scenarios. The benefits of this approach are demonstrated in Fig. 7, which presents simulated future distributions for *R. alba* generated by passing climate inputs generated from the 2050s high scenario of Hulme and Jenkins (1998) through ANNs trained using both the hierarchical and non-hierarchical approaches. Land-cover was kept constant for this example. Simulations using both models show the expected general trend of increasing suitability as the species climate envelope moves northward (cf. Fig. 2). However, there are notable differences between the predictions, particularly in the south west of England where the non-hierarchical approach shows a marked increase in suitability. Examination of the observed European species distribution and of the European-scale model inputs aids interpretation of these results. It is plausible that the southern range margin of *R. alba*, which is currently in southern France and northern Spain, will limit the species distribution in south west England under the climate change scenario for the 2050s. Whilst the hierarchically-trained model identified this trend, keeping suitability relatively low, the non-hierarchical model was forced to predict for climatic regimes for which it had no parameterisation, leading to an arbitrary simulation of increased suitability in the south west.

A number of limitations to the modelling approach have been noted, including the restricted explanatory power of ANNs and the reliance on correlations between observed distributions and environmental variables. Further limitations are inherent in the availability and accuracy of datasets. Data can rarely be generated for all resolutions and for all spatial extents, but rather tends to be available for large extents at coarse resolutions, or small extents at fine resolutions. Thus, in this study species distributions were obtained at 0.5° resolution for Europe, 10 km resolution for Britain, and 2 km resolution for local case studies. It has been necessary to design the modelling framework to take best advantage of the available data. Questions regarding the accuracy of the data also arise, in particular regarding the assumption that observed species absences are true absences, and not a result of insufficient sampling (Griffiths et al. 1999). The use of species records spanning many years is also potentially problematic, since distributions are dynamic over relatively short time-scales. The use of mean 1961–1990 climate data aims to reduce this effect, though the year 2000 ‘snapshot’ of land-cover will add an element of error to the modelling. Base errors arising from data limitations are unavoidable. However, the level of success that has been

achieved in modelling species distributions has demonstrated that biogeographical trends can be identified regardless of the imperfect data that is so often all that is available in ecological studies.

In conclusion, this paper has presented a novel modelling methodology for integrating climate and land-cover variables across different spatial scales. Application of the model to four plant species in Britain has enabled the methodology to be tested and conclusions to be drawn regarding the relative roles of climate and land-cover in influencing distributions. Further application of the approach to other species and landscapes should help further our understanding of the potential combined effects of climate and land-cover change on the distribution of species.

Acknowledgements – We are very grateful to P. Berry, R. Økland, R. Smithers, W. Thuiller and M. Trivedi for useful advice and comments on the manuscript. We thank I. Brown and M. Colley (U.K. Climate Impacts Programme) for help obtaining climate and land-cover data, and H. Arnold (U.K. Biological Records Centre) and M. Sanford (Suffolk Biological Records Centre) for providing species distribution data. The research was funded by the European Community's Fifth Framework Programme (ACCELERATES project, contract EVK2-CT-2000-000610) and by the MONARCH project, which was funded by English Nature (lead), Countryside Council for Wales, Environment Agency, Environment and Heritage Service, Forestry Commission, Joint Nature Conservation Committee, National Parks and Wildlife (Republic of Ireland), National Trust, Royal Society for the Protection Birds, Scottish Executive, Scottish Natural Heritage, U.K. Climate Impacts Programme, Welsh Assembly Government and Woodland Trust. We thank these organisations for their intellectual input and practical support.

References

- Beckett, G., Bull, A. and Stevenson, R. 1999. A Flora of Norfolk. – Gillian Beckett.
- Berry, P. M. et al. 2002. Modelling potential impacts of climate change on the bioclimatic envelope of species in Britain and Ireland. – *Glob. Ecol. Biogeogr.* 11: 453–462.
- Berry, P. M. et al. 2003. The sensitivity and vulnerability of terrestrial habitats and species in Britain and Ireland to climate change. – *J. Nat. Conserv.* 11: 15–24.
- Cohen, J. 1960. A coefficient of agreement for nominal scales. – *Educ. Psychol. Meas.* 20: 37–46.
- Collingham, Y. C., Hill, M. O. and Huntley, B. 1996. The migration of sessile organisms: a simulation model with measurable parameters. – *J. Veg. Sci.* 7: 831–846.
- Collingham, Y. C. et al. 2000. Predicting the spatial distribution of non-indigenous riparian weeds: issues of spatial scale and extent. – *J. Appl. Ecol.* 37: 13–27.
- Fielding, A. H. and Bell, J. F. 1997. A review of methods for the assessment of prediction errors in conservation presence/absence models. – *Environ. Conserv.* 24: 38–49.
- Franklin, J. 1995. Predictive vegetation mapping: geographic modelling of biospatial patterns in relation to environmental gradients. – *Prog. Phys. Geogr.* 19: 474–499.
- Gevrey, M., Ioannis, D. and Lek, S. 2003. Review and comparison of methods to study the contribution of variables in artificial neural network models. – *Ecol. Modell.* 160: 249–264.
- Griffiths, G. H., Eversham, B. C. and Roy, D. B. 1999. Integrating species and habitat data for nature conservation in Great Britain: data sources and methods. – *Glob. Ecol. Biogeogr.* 8: 329–345.
- Guisan, A. and Zimmermann, N. E. 2000. Predictive habitat distribution models in ecology. – *Ecol. Modell.* 135: 147–186.
- Halliday, G. et al. 1997. A Flora of Cumbria: comprising the vice-counties of Westmorland with Furness (v.c.69), Cumberland (v.c.70) and parts of north-west Yorkshire (v.c.65) and north Lancashire (v.c.60). – Lancaster Univ.
- Hanley, J. A. and McNeil, B. J. 1982. The meaning and use of the area under a Receiver Operating Characteristic (ROC) curve. – *Radiology* 143: 29–36.
- Hilbert, D. W. and Ostendorf, B. 2001. The utility of artificial neural networks for modelling the distribution of vegetation in past, present and future climates. – *Ecol. Modell.* 146: 311–327.
- Hulme, M. and Jenkins, G. J. 1998. Climate change scenarios for the United Kingdom: scientific report. – U.K. Climate Impacts Programme Technical Report No.1, Climatic Research Unit, Norwich.
- Hulme, M. et al. 1995. Construction of a 1961–1990 European climatology for climate change modelling and impact applications. – *Int. J. Climatol.* 15: 1333–1363.
- Ivimey-Cook, R. B. 1984. Atlas of the Devon Flora: flowering plants and ferns. – Devonshire Ass. Adv. Sci., Lit. Arts.
- Lek, S. and Guegan, J. F. 1999. Artificial neural networks as a tool in ecological modelling: an introduction. – *Ecol. Modell.* 120: 65–73.
- Maier, H. R. and Dandy, G. C. 1998. Understanding the behaviour and optimising the performance of back-propagation neural networks: an empirical study. – *Environ. Modell. Softw.* 13: 179–191.
- Manel, S., Williams, H. C. and Ormerod, S. J. 2001. Evaluating presences-absence models in ecology: the need to account for prevalence. – *J. Appl. Ecol.* 38: 921–931.
- Martinez, J. A., Serrano, D. and Zuberogoitia, I. 2003. Predictive models of habitat preferences for the Eurasian eagle owl *Bubo bubo*: a multi-scale approach. – *Ecography* 26: 21–28.
- May, R. M. 1989. Levels of organisation in ecology. – In: Cherratt, J. M. (ed.), *Ecological concepts*. Blackwell, pp. 339–363.
- Monserud, R. A. and Leemans, R. 1992. Comparing global vegetation maps with the Kappa statistic. – *Ecol. Modell.* 62: 275–293.
- Openshaw, S. 1984. The modifiable areal unit problem. – *Geo Books*.
- Parmesan, C. and Yohe, G. 2003. A globally coherent fingerprint of climate change impacts across natural systems. – *Nature* 421: 37–42.
- Pearce, J. and Ferrier, S. 2000. Evaluating the predictive performance of habitat models developed using logistic regression. – *Ecol. Modell.* 133: 225–245.
- Pearson, R. G. and Dawson, T. P. 2003. Predicting the impacts of climate change on the distribution of species: are bioclimate envelope models useful? – *Glob. Ecol. Biogeogr.* 12: 361–371.
- Pearson, R. G. et al. 2002. SPECIES: a spatial evaluation of climate impact on the envelope of species. – *Ecol. Modell.* 154: 289–300.
- Perring, F. H. and Walters, S. M. (eds) 1962. Atlas of the British Flora. – Thomas Nelson and Sons.
- Preston, C. D., Pearman, D. A. and Dines, T. D. (eds) 2002. New Atlas of the British and Irish Flora. – Oxford Univ. Press.
- Rettie, W. J. and Messier, F. 2000. Hierarchical habitat selection by woodland caribou: its relationship to limiting factors. – *Ecography* 23: 466–478.
- Sala, O. E. et al. 2000. Global biodiversity scenarios for the year 2100. – *Science* 287: 1770–1774.

- Swets, J. A. 1988. Measuring the accuracy of diagnostic systems. – *Science* 240: 1285–1293.
- Tarassenko, L. 1998. A guide to neural computing applications. – Arnold.
- Thuiller, W. et al. 2003. Large-scale environmental correlates of forest tree distributions in Catalonia (NE Spain). – *Glob. Ecol. Biogeogr.* 12: 313–325.
- Travis, J. M. J. 2003. Climate change and habitat destruction: a deadly anthropogenic cocktail. – *Proc. R. Soc. Lond. B Biol.* 270: 467–473.
- Warren, M. S. et al. 2001. Rapid responses of British butterflies to opposing forces of climate and habitat change. – *Nature* 414: 65–69.