Predicting vegetation condition and weed distributions for systematic conservation management

An application of GRASP in the central South Island, New Zealand

SCIENCE FOR CONSERVATION 220

Jacob McC. Overton and Anthony Lehmann

Published by
Department of Conservation
PO Box 10-420
Wellington, New Zealand

Science for Conservation is a scientific monograph series presenting research funded by New Zealand Department of Conservation (DOC). Manuscripts are internally and externally peer-reviewed; resulting publications are considered part of the formal international scientific literature.
Titles are listed in the DOC Science Publishing catalogue on the departmental website http://www.doc.govt.nz and printed copies can be purchased from science.publications@doc.govt.nz

ISSN 1173–2946
ISBN 0–478–22403–6

In the interest of forest conservation, DOC Science Publishing supports paperless electronic publishing. When printing, recycled paper is used wherever possible.

This report was prepared for publication by DOC Science Publishing, Science & Research Unit; editing and layout by Lynette Clelland. Publication was approved by the Manager, Science & Research Unit, Science Technology and Information Services, Department of Conservation, Wellington.

CONTENTS

Predicting vegetation condition and weed distributions for systematic conservation management

An application of GRASP in the central South Island, New Zealand

Jacob McC. Overton[1] and Anthony Lehmann[2]

[1] Landcare Research, Private Bag 3127, Hamilton, New Zealand
[2] Swiss Center for the Cartography of Fauna, Terreaux 14, CH-2000 Neuchâtel, Switzerland

A B S T R A C T

The New Zealand Department of Conservation's Measuring Conservation Achievement (MCA) process is an ambitious and comprehensive attempt to inform conservation decision-making and to report on the difference made by conservation management actions using a rigorous foundation of research and empirical data. Generalised regression analysis and spatial prediction (GRASP) provides several key functions for MCA, and indicator programmes in general, including: 1) spatially explicit estimates produced by spatial extrapolations from point data; 2) regional descriptions of biodiversity patterns and factors influencing biodiversity (or other characteristics of interest); and 3) methods to identify and correct biases in data that result from ad-hoc or biased sampling patterns. The use of GRASP to provide information for MCA is developed and demonstrated, using vegetation data from the central South Island of New Zealand. Vegetation condition indices and weed distributions were combined with climatic, landform, land cover, cadastral and fragmentation variables to produce spatial predictions of the biodiversity attributes measured on the plots. A sampling bias towards vegetation in better condition, resulting from biased plot distribution, was corrected using GRASP. Vegetation condition indices were successfully modelled and showed different patterns on conservation and non-conservation lands. Weed distributions were successfully modelled with composite models, and showed a wide variation between species in their relationships to environmental and other variables. Together, these provide spatially explicit predictions of condition and pressure for use in conservation planning and reporting. The results demonstrate the utility of GRASP for systematic conservation management, and highlight the need to invest in underlying data collected according to well-defined sampling schemes.

Keywords: Measuring conservation achievement, management, generalised regression analysis and spatial prediction, geographic information systems, biodiversity, weeds, vegetation condition, probability sampling, found data, New Zealand

Overton, J. McC.; Lehmann, A. 2003: Predicting vegetation condition and weed distributions for systematic conservation management: an application of GRASP in the central South Island, New Zealand. *Science for Conservation 220*. 57 p.

1. Introduction and background

The 1990s to early 2000s has seen a developing trend towards conservation management becoming more systematic and rigorous. A large part of this trend has focused on prioritising and optimising reserve design and selection (e.g. Margules & Pressey 2000). Underlying this are advances in the prediction and depiction of the spatial distribution[1] of conservation assets such as species (e.g. Franklin 1995; Cawsey et al. 2002; Ferrier et al. 2002a) or communities (Ferrier et al. 2002b). Conservation management is, of course, about more than just reserve design, and to fully develop these systematic approaches requires that systematic conservation planning be generalised beyond the design of reserve systems to incorporate planning for the entire range of conservation activities, as well as reviewing and reporting on these activities.

Most conservation agencies must juggle a diverse range of activities, all competing for the same limited budget. This situation is particularly pronounced in New Zealand (and many other island ecosystems) where biodiversity conservation is not only about mitigating direct human impacts, such as logging, land clearance and hunting, (which are generally excluded from reserve areas) but also indirect impacts, such as introduced predators, pests, and weeds. Thus, information is required not only on the spatial distribution of assets, but also on the condition of the assets, and the pressures on biodiversity, such as the prediction of alien plants (Higgins et al. 1999) and the dynamics of invasive pines (Higgins et al. 1998).

Systematic conservation management can be viewed as the extension of systematic conservation planning beyond reserve design to the implementation of the full range of conservation activities and the review and reporting of these activities against conservation objectives. Methods and approaches are required for explicit and systematic planning and prioritising between different sorts of places (e.g. forest reserves, marine reserves), and between very different sorts of activities (e.g. weed control, buying new reserves, research, public education, improving visitor facilities). Methods and approaches are also developed for the systematic review of current or past actions, and reporting on outcomes achieved and resources required for achieving goals. Systematic conservation management can be viewed as an integrated framework for planning, implementation, reporting and review of conservation actions.

The Measuring Conservation Achievement process (MCA, Stephens 1999; Stephens et al. 2002) developed by the New Zealand Department of Conservation (DOC) is a comprehensive and ambitious attempt to inform national conservation decisions and report on the difference made by conservation activities using an objective, rigorous basis. The power of MCA is that it provides a framework for systematic conservation management by informing allocation decisions and allowing reporting on a broad range of conservation assets and activities. However, this power comes at a cost: the MCA process is information hungry; it requires spatially explicit information

[1] Many of the terms used in this report are explained in a glossary, see p. 55.

about a multitude of diverse conservation threats and assets. The broad-scale, comprehensive and quantitative approach used by MCA is made feasible only by the recent and continuing exponential growth in computing power and information processing capability of computers.

With the MCA process, DOC has the opportunity to produce a world first: conservation decision making and reporting with an explicit and largely quantitative basis. However, this opportunity can only be realised with investment in the information foundation for MCA. To be effective, MCA needs a strong foundation in data and research. Reams of raw data are not useful for MCA; it requires derived measures arising from generalised and integrated information. This process of integrating different sorts of information and generalising it into forms useful for decision making can be represented by an information pyramid (Fig. 1). This process provides an excellent example of informed conservation decision making, as described by Overton et al. 2002. Without investment in this foundation, DOC will be in a situation all too common within conservation management agencies—that of making decisions based on little or no data, without objective methods for integrating the information already available, and being quite unable to report on what difference conservation effort is making to the state of natural heritage (Fig. 2).

1.1 GENERALISED REGRESSION ANALYSIS AND SPATIAL PREDICTION (GRASP)

GRASP (Lehmann et al. 2002) is both a general concept and a specific implementation in Splus (Chambers & Hastie 1993). GRASP combines the power of advanced statistical and spatial analyses with the advantages of extensive spatial information managed in Geographic Information Systems (GIS).

The GRASP process is perhaps best understood as a means of defining the patterns of the response variable in relation to climatic, landform or other spatial predictor variables and using these patterns to make the best prediction of the response variable across the landscape. This process can be illustrated with a single response variable and a single spatial predictor variable. Consider the case of using a GIS surface of mean annual temperature to make spatial extrapolations of the abundance of a weed species for a region. In this case, GRASP is a simple regression technique, using plot observations of weed abundance to produce a regression model of weed abundance as a function of temperature. Using this model and the surface of mean annual temperature, one could then predict weed abundance for the entire region.

The GRASP process can also use categorical predictor variables, such as land cover or soil type. Consider a model of weed abundance as a function of land cover. In this case, the GRASP model would just be an analysis of variance, consisting of the mean weed abundance for each land cover category. The prediction of weed abundance across the region would be simply the average weed abundance for each land cover class, applied to all places with this particular land cover.

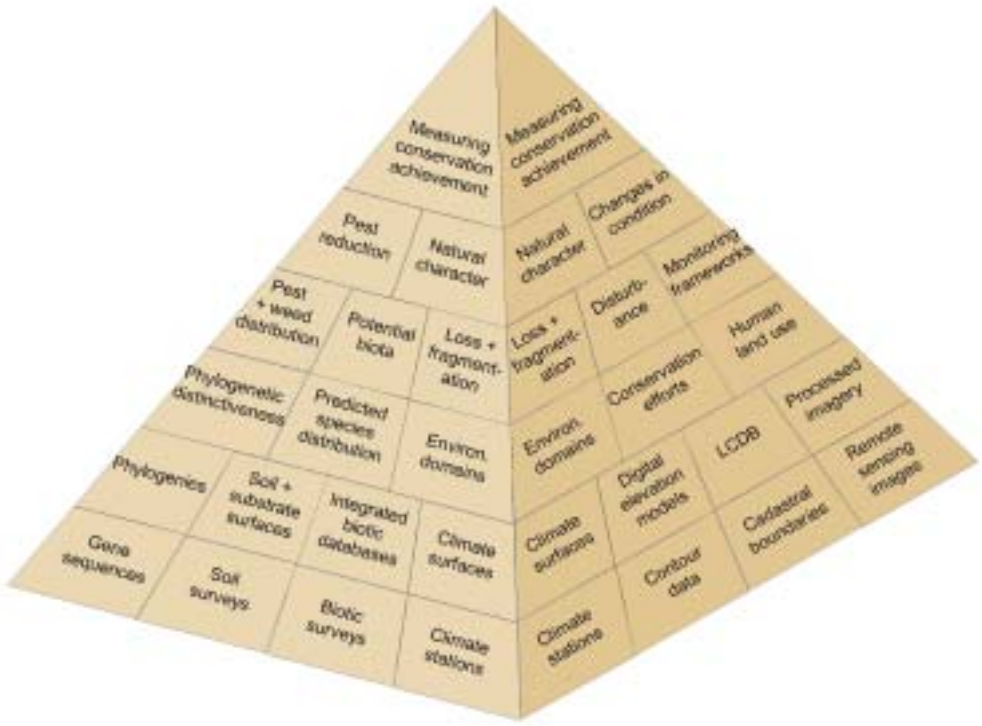

Figure 1. Information pyramid, showing MCA at the apex. Here, the MCA process is supported by quantitative data and rigorous methods for integrating and generalising the information. After Overton et al. (2002).

Most models in GRASP are more complex than either of these examples, and many include a combination of continuous variables and categorical variables. The overall approach, however, remains the same. Essentially, GRASP takes each variable of interest and tries to find the pattern of this variable across the landscape, relative to environmental, land cover, fragmentation, or spectral characteristics. It then uses this pattern to predict the variable of interest for the entire region.

More generally, GRASP can be seen as a combination of regression modelling and spatial prediction. Modern regression analysis is used to establish relationships between a response variable and a set of spatial predictors. The regression models are then used to make spatial predictions of the response. The GRASP process requires point measurements of the response, as well as regional coverages of predictor variables that are statistically (and functionally) important in determining the patterns of the response. In the study reported here, point measurements of vegetation condition and weed abundance (the response variables, used one at a time) were regressed against environmental variables such as climate, landform, land cover and fragmentation (the spatial predictors). These regression relationships can then be used to predict the species abundance from the environmental surfaces.

The GRASP implementation and graphical user interface (GUI) are designed to facilitate the GRASP process and the analyses needed to check the models and predictions. The implementation consists of a collection of functions and user interface in Splus that provide a toolbox for quick and easy data checking, model building and evaluation, and calculation of predictions. In addition to making the GRASP process easier, this implementation also standardises the

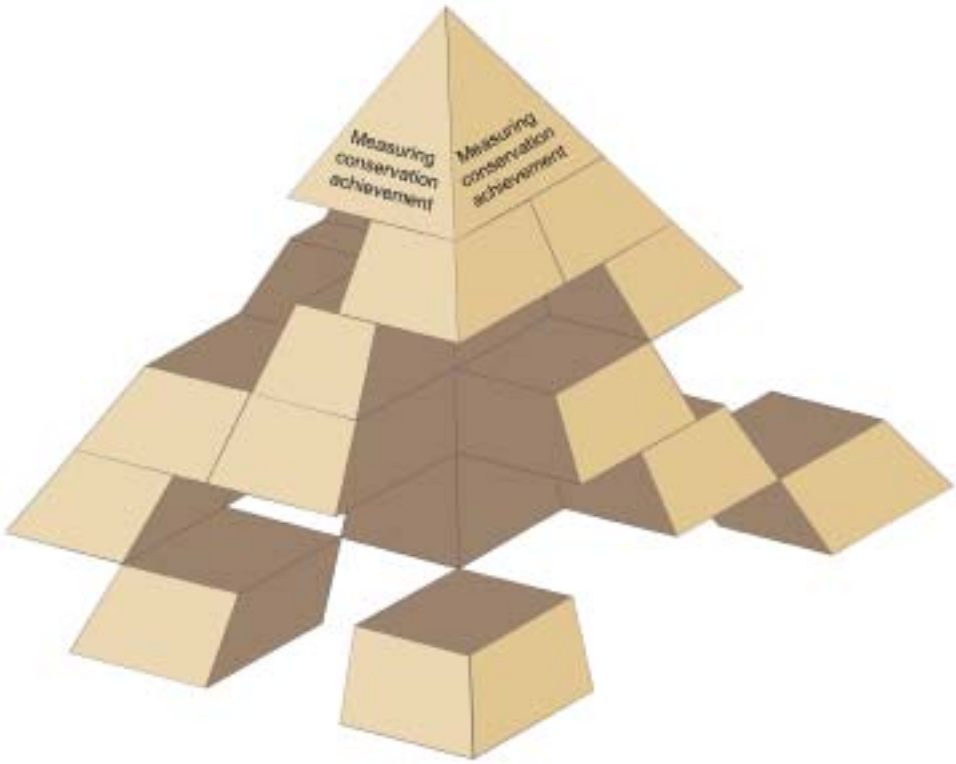

Figure 2. Information pyramid without a foundation. This is the scenario resulting from a lack of investment in research and monitoring and methods for integrating and generalising the information.

modelling process and makes it more reproducible and less subjective, while preserving flexibility. The current version uses generalised additive models (GAMs), a modern non-parametric regression technique with a number of advantages for ecological modelling. GRASP is a leading technique for making these predictions, and has been developed in conjunction with, and been reviewed by, international leaders in the field, many of whom actually use GRASP in their work. However, GRASP is still relatively new, and refinements and improvements are still being developed.

1.1.1 GRASP and monitoring programmes

GRASP performs three powerful roles in MCA, and indicator programmes in general. Each of these roles is demonstrated in this report.

1. GRASP provides robust spatial predictions (maps) across entire regions from plot-based measurements. By using the environmental information to make spatial extrapolations from plot measurements, GRASP provides the best estimate for the spatial pattern of the indicator across the region. This is illustrated in this report by the maps of the vegetation condition indices and weed abundances.
2. GRASP provides methods for dealing with the many problems associated with non-representative (biased) data from past surveys or data collected as a by-catch from other conservation work. The regression techniques used in GRASP offer a number of opportunities to account for a biased distribution of plots across the landscape. This correction is inherent in the regression approach, and can be enhanced by estimating plot densities, as is done in this report.

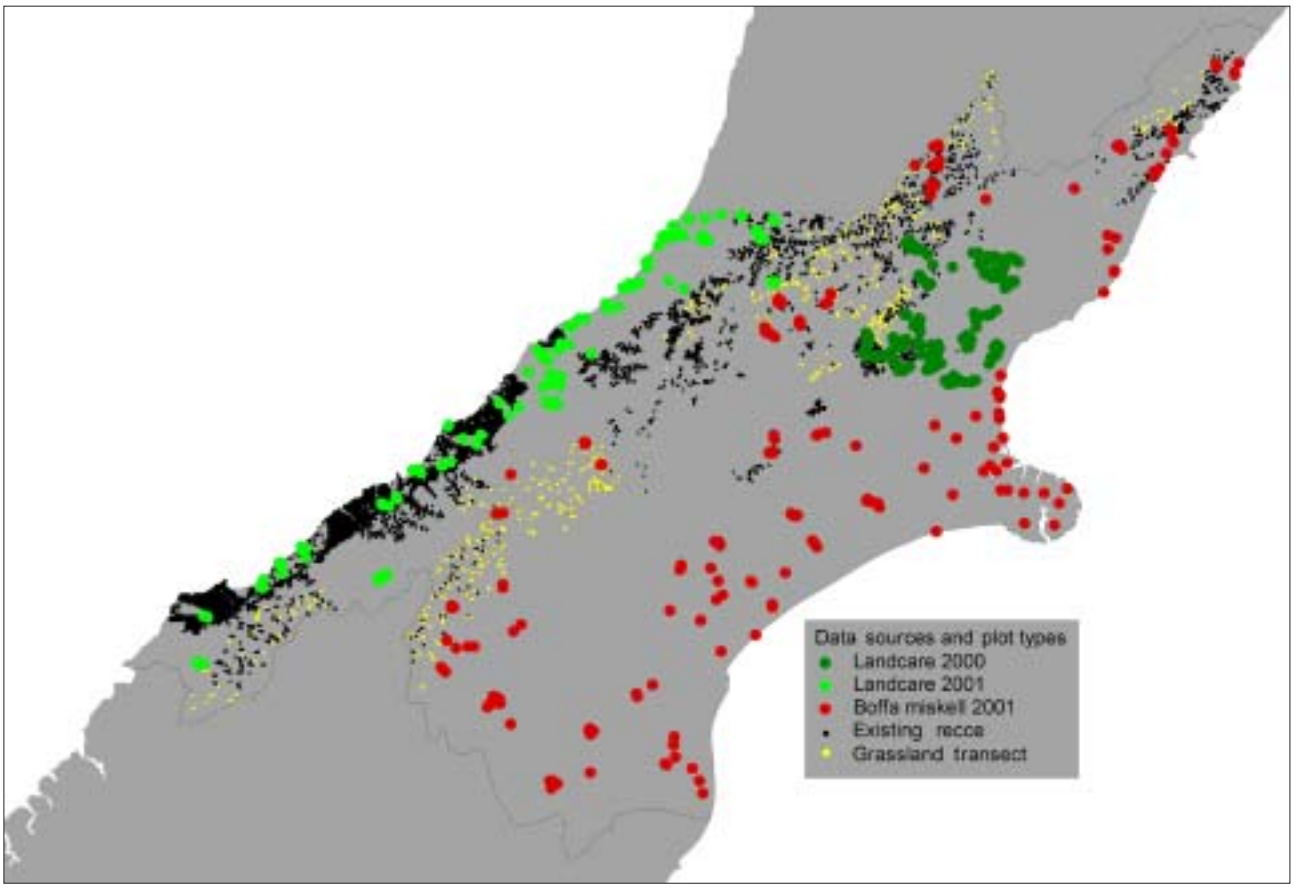

Figure 3. Locations of samples used in the analyses. The different data sources are shown, along with the boundary of the study region.

3. GRASP develops the regional relationships between indicators and environmental factors. These relationships are important for understanding biodiversity patterns and influences on them at regional scales, and to provide insight for management and research. In this report, the regression models show relationships between the vegetation condition indices and weed abundances (response variables) with the environmental variables such as climate, substrate, land use, or fragmentation (predictor variables).

1.2 SAMPLING METHODS AND FOUND DATA

1.2.1 Probability sampling

Consider the problem faced in this study. We are attempting to estimate vegetation condition and weed distributions for a study region of almost 6 million hectares (Fig. 3). Clearly, it is impractical to measure every hectare of the region. Methods are required to sample the area and to make estimates for the whole region from this sample. While much of statistics (including the methods in GRASP) deals with what to do once you have the data, sampling methods deal with where and how to gather the data. Probability sampling refers to sampling designs in which the locations or elements to sample are chosen according to a random sampling scheme. Examples include simple random sampling, stratified random sampling and most grid sampling.

Correct sampling methods are important for ensuring that samples are not biased and accurately reflect the population of interest. Without a sound sampling procedure, or appropriate methods for dealing with the found data (see below), there is no assurance that the sampled measurements are an accurate representation of the true population. One of the fundamental requirements for a sampling scheme is to define the population that we want to characterise. For this study, we can define our population of interest as every hectare of the study region. More generally for MCA, we could define it as every hectare of New Zealand.

Two recent examples exist of the use of probability sampling of ecosystem characteristics in New Zealand. Hall et al. (2001) and Coomes et al. (2001) used a combination of existing plots and new plots on a grid sampling scheme to establish baseline and monitoring schemes for carbon storage in indigenous forests. Overton et al. (2000, 2002) used probability sampling schemes and GRASP to characterise roadside biodiversity for the Waikato Region.

Another component of sampling design is selecting which plot methods to use. This study used standard recce (see Section 3.2.4) plots to characterise vegetation. Since existing plot methods were designed for different purposes, it may well be worthwhile to develop new plot methods for characterising biodiversity, but this topic is not dealt with in this report.

1.2.2 Sampling design

General sampling schemes, such as systematic grid sampling, are excellent ways to achieve an unbiased characterisation of a population. Using grid sampling or simple random sampling, the samples are distributed across subpopulations in proportion to the abundance of each in the overall population.

Sometimes, however, certain subpopulations may be of much more interest than others. As a result, it may be more efficient (i.e. more information can be gained for the same money) to sample more heavily in areas of greater interest. This is called stratified sampling. A recent example of stratified sampling of biodiversity in New Zealand is the roadside biodiversity sampling of Overton et al. (2002). It was known from previous work (Overton et al. 2000) that roadside indigenous biodiversity was much higher when the road passed through forest than when it passed through pasture. Yet only 5% of the roadsides were next to forest, while 70% were next to pasture. A systematic sample of 300 sites would have achieved about 210 pasture sites but only 15 forest sites. In this example sampling was stratified by land cover, and roadsides through forests were sampled more intensively than roadsides through pasture to achieve a roughly equal number of pasture and forest sites.

Stratified sampling essentially introduces bias into the sampling by sampling some subpopulations more heavily than others. However, the bias can be corrected by using the sampling intensities to weight the analyses and remove the bias. The result is unbiased estimates of the overall population with more information for a given amount of money.

1.2.3 Found data

Found data has been defined by Overton et al. (1993) as data for which the sampling design is unknown. This may be because the data was not collected using a specific design, or because the design has been lost, or because the original design reflected other constraints or objectives. The main difficulty with found data is that the sampling intensities are not available for use in analysis. For instance, if the data came from a biased design, then to achieve an unbiased result, the sampling intensities would need to be taken into account in the analyses. However, sometimes sampling intensities are not available. Since sampling design specifications are one of the most important pieces of information that must be known for the rigorous use of data, sampling design specifications should be a standard part of metadata.

The existing data used in this study are found data. They consist of a collection of surveys done at different times and places, with different sampling designs. It is also clear that the existing data are biased towards particular parts of the landscape. Certain regions and certain land covers, like indigenous forests, are much more heavily sampled than, for example, pastoral farming land in the Canterbury plains. If one is interested in getting the overall value of a vegetation condition index for the study region, it is clear that a simple mean of the existing data will not suffice. Since the existing data are biased towards indigenous forest in good condition, this would result in an erroneously high estimate of vegetation condition across the region.

Overton et al. (1993) proposed two solutions to the problem of found data. One consisted of estimating the inclusion probabilities, and the other of using regression techniques to account for the bias. GRASP is very useful for both of these approaches, and both are used in this study.

Found data are not free data. A large amount of work is required to use such data in a robust fashion. However, it should be noted that one of the most valuable uses of past data is in the characterisation of historical patterns, such as developing 1990 baselines for carbon monitoring (Hall et al. 2001), or characterising natural distributions of species in undisturbed forests (e.g., Leathwick 1998, 2001; Leathwick & Austin 2001; Leathwick & Whitehead 2001; Lehmann et al. 2002a), or past patterns of weed invasion (Wiser et al. 1998).

2. Objectives

The objectives of this study were to:

1. Develop and demonstrate the use of GRASP to provide for the information needs of the MCA process, including the production of spatial predictions, regional patterns of biodiversity, and methods to deal with biased data.
2. Provide predictions of vegetation condition and weed distributions for the implementation of MCA in the Canterbury Region.
3. Identify the properties of underlying data, including sampling design, needed to provide reliable predictions.
4. Investigate the characterisation of fragmentation, and the influence of fragmentation on regional patterns of biodiversity.

3. Methods

3.1 STUDY AREA

The study region was defined as the entire Canterbury Region, and the portion of the West Coast Region south of 42°37′51″ S (Fig. 3). The study region includes a wide range of climatic and environmental characteristics, and a correspondingly wide range of ecosystem characteristics and land uses. This region encompasses some of the best protected ecosystems of New Zealand (e.g. the high Southern Alps), as well as the most poorly protected (e.g. the Canterbury Plains).

3.2 VEGETATION DATA

While a broad range of existing vegetation surveys and plots were available in the National Vegetation Survey (NVS) databank, only a subset was used in these analyses. New surveys were designed to fill in gaps of environmental, land-cover and land-tenure combinations that were either unrepresented or under-represented in the existing data.

3.2.1 Landcare Research Canterbury surveys 2000

These vegetation surveys were conducted by Landcare Research in January and February 2000 in a region defined by a SPOT4 satellite image. Surveys were designed to cover a range of environmental and spectral (i.e. reflectance from SPOT4 imagery) characteristics. While these surveys covered a wide gamut of land cover and vegetation types, in both crown and private lands, they were limited to a 60 × 60 km square in the Canterbury plains, and thus covered only a limited part of the study area.

Plots were bounded 20 × 20 m recce plots with cover recorded directly in cover percentages (not Braun-Blaunquet categories). For 1 out of 10 plots, point height frequency methods were done on the same plots to investigate their utility in calibrating cover estimates between observers and across cover types. On a grid with 2-m spacing, 25 points were chosen within the plot. For each point on the grid, a 10 × 10 cm subquadrat was placed over the point, and any vegetation that entered the column above this subquadrat was recorded by species and height of intersection.

3.2.2 Landcare Research West Coast surveys 2001

These surveys were designed specifically for this project to provide additional information on vegetation in the portion of the study area in the West Coast Region. The purpose of the additional sampling was to gather more data in environmental and land-cover combinations poorly represented with little existing data, as well as to provide some new information in the more heavily sampled areas. Sample locations were chosen subjectively, using available GIS information on land cover, environmental domains, DOC boundaries, and existing plot locations to position new plots. The new sample locations were then exported as eastings and northings of the New Zealand map grid (NZMG) and uploaded to Garmin etrex Global Positioning System (GPS) units. Field maps were produced in a GIS by overlaying plot locations onto digital copies of topographic maps. Field crews used the field maps to navigate near the plots and the GPS to find the final location for sampling. The chosen plot location was always the southwest corner of the plot, with the plot extending 20 m north and 20 m east of this origin. In cases where the GPS was not giving readings, azimuths and directions to plots were calculated from a nearby landmark, and plot locations were estimated to the best possible accuracy.

Plots were bounded 20 × 20 m recce plots with cover recorded directly in cover percentages (not Braun-Blaunquet categories). For 1 out of 5 of the plots, the calibration plots of Landcare Research Canterbury Surveys 2000 were used. This calibration methodology was not part of the contracted work, and the analyses are not reported here, but these methods form a potentially valuable method for investigating the subjectivity of cover estimates.

3.2.3 Boffa Miskell Canterbury surveys 2001

These surveys were also designed to increase information in environmental and land-cover combinations poorly represented by existing data, as well as to provide some new information in the more heavily sampled areas. The surveys were performed by a number of staff of Boffa Miskell Consulting, under contract to the Canterbury Conservancy of DOC. Because of political considerations, no sampling was done on non-DOC land in the Canterbury Region. Plot methods used were bounded 20 × 20 m recce plots with cover recorded directly in cover percentages (not Braun-Blaunquet categories).

3.2.4 NVS data

Available plot locations were sourced from the NVS database (Wiser et al. 2001), selecting recce plots and grassland transects only and excluding any of these that were considered to be not representative of the overall vegetation in the area, such as Protected Natural Area (PNA) surveys or exclosure plots.

Recce plots: Recce plots are either bounded 20 × 20 m plots or unbounded plots of approximately that dimension. For some of the past recce surveys, cover scores were recorded for (usually) six different height tiers. Other existing recce surveys recorded only species presence (and dominance) in the tiers. See Allen (1992) for explanation of Recce plots and methods.

Recce plots with cover scores for tiers were used in modelling variables that require cover, including the vegetation indices that require cover listed in Table 1, and the weed abundance models. Recce plots with presence in tiers were only used to estimate variables that require species presence, including vegetation indices that do not require cover and weed presence-absence models.

Grassland transects: Grassland transect data consisted of the summarised data from the central transect of grassland plots. The summarised data give the frequency of each observed species for the (usually 50) circular subquadrats along the central transect. See Wiser & Rose (1997) for details of grassland transect methods.

While the species frequencies of grassland transects give numbers similar to estimated cover, these numbers are not directly comparable. For instance, for grassland transects, sums over all species of species frequencies (expressed as percentages) often exceeded 500, while for recce plots, sums of percent cover over species and tiers rarely exceeded 300. Therefore, grassland transects were only used to provide species presence-absence, and were used to model the same variables as recce plots with presence only (i.e. the vegetation indices that do not require cover and the weed presence-absence models). It should be noted that the presence and absence of species derived from the grassland transects is also not directly comparable to that derived from the recce plots because of the different plot methods and different area sampled. But because there was so little other data available in grasslands, the use of grassland transect data for these purposes was considered justified.

3.3 VEGETATION CONDITION INDICES

A number of different vegetation condition indices were derived from the plot data and modeled using GRASP. Table 1 defines the variables and types of models used.

3.4 WEED SPECIES

A number of weed species of conservation concern were considered for modelling. An original list of 15 species was developed in consultation with Canterbury and West Coast Conservancies. Species were chosen to be of specific conservation concern, but also present in the data. In some cases, groups of species—such as willows or wilding conifers—were lumped because of similarities amongst the species with respect to conservation problems and control action. An additional 30 species were chosen as candidates to model, based on general interest and availability of data. Seven of these species were eliminated because of excessive records of unidentified congeners or other uncertainties in identification (weed species are listed later in Table 4).

TABLE 1. VARIABLES MODELED IN THIS STUDY.

Vegetation condition indices are defined and method of calculation shown. Weed models are composite models of presence-absence models and abundance models. See Table 4 for detailed information on each weed species. For each modelled variable, the type of model and subset of plots used to model them is shown, with sample sizes.

NAME	ABBREVIATION	DEFINITION	CALCULATION	UNITS	TYPE OF MODEL	PLOT SUBSET USED IN MODELS
Native species richness	NatRich	Number of native species	S_n	Count	Poisson	All plots, n = 16672
Exotic species richness	ExotRich	Number of exotic species	S_e	Count	Poisson	All plots, n = 16672
Proportion of species native	PSppNat	The proportion of species in the plot that are native	S_n/S	Proportion	Binomial	All plots with S > 0, n = 16640
Cover	Cover	Total percentage cover for plot, summed across all species and tiers	$\Sigma_{spp}\Sigma_{tiers}C_{ij}$	$m^2/100\ m^2$	Poisson	Recce plots with cover recorded, n = 9351
Proportion cover native	PCovNat	The proportion of cover accounted for by native species	(Native Cover)/Cover	Proportion	Binomial	Recce plots with cover recorded and S > 0, n = 9319
Biomass index[1]	Biomass	The sum, across species and tiers, of cover × height of the tier	$\Sigma_{spp}\Sigma_{tiers}C_{ij}h_j$	$m^3/100\ m^2$	Poisson	Recce plots with cover, n = 9351
Proportion biomass native	PBioNat	The proportion of the biomass index accounted for by native species	(Native Biomass)/Biomass	Proportion	Binomial	Recce plots with cover recorded and S > 0, n = 9319
Weed presence-absence	SppcodPA[2]	The probability that a species occurs in a plot	NA	Probability	Binomial	All plots, n = 16672
Weed abundance given presence	SppcodAB[2]	The total cover of a species where it is present	$\Sigma_{tiers}\Sigma_j$	$m^2/100\ m^2$	Poisson	Recce plots with cover recorded and species present[3]
Weed abundance composite	none	The estimated abundance for a species	SppcodPA*SppcodAB	$m^2/100\ m^2$	Not modeled directly	Not modeled directly

[1] As calculated, this might be better named Vegetation Volume

[2] Sppcod is the six letter species code for each weed species, see Table 4 for list of species

[3] See Table 4 for numbers of plots with each species

S_n = Number of native species in plot

S_e = Number of exotic species in plot

S = Total number of species in plot (species richness)

C_{ij} = Percent cover of species i in tier j

h_j = height of tier j

3.5 SPATIAL PREDICTORS

Spatial predictors developed for this study included a number of climatic, landform and land cover variables, as well as fragmentation and cadastral variables (Table 2). Fragmentation variables can be grouped into two types—patch-based and continuous. Both types were developed for this study. Many of the spatial predictors developed were not used in modeling, for the following reasons.

A number of variables were excluded from analyses because they were highly correlated with other variables. A correlation matrix of all spatial predictor variables was produced, and variables were removed so that no variable pairs had correlations above 0.85. Excluded variables included a number of climatic variables, and fragmentation variables. Surprisingly, DOC300 and DOC1000 had a fairly low correlation (see Table 2).

There were a number of problems with substrate variables derived from the New Zealand Land Resource Inventory (NZLRI) that led to them being excluded from models. These problems included: (1) incomplete coverage for the study area (which would have resulted in incomplete predictions for any model using those variables), and (2) poor mapping accuracy that produced spatial artifacts at fine spatial scales. Furthermore, much of the relevant fine spatial pattern in substrate was reflected in differences in land cover, which was provided by the Land Cover Database (LCDB), which provides polygon predictions of basic land cover categories, produced by a combination of automatic and manual interpretation of SPOT satellite imagery, with a minimum mapping unit of one hectare.

Many of the patch-based fragmentation variables had statistical distributions that made them difficult to model, and a number of operational and conceptual difficulties (see Section 5.5) that led to them not being included in the models.

Unfortunately, consistent remote sensing imagery, such as SPOT4 or LANDSAT images, was not available for the entire study area, so could not be used as a spatial predictor. However, a SPOT4 image for a portion of the study region was used to explore the advantages of remote sensing information for this work.

3.6 GRASP

Generalised Regression Analysis and Spatial Prediction (GRASP, Lehmann et al. 2002b) was used to model the variables shown in Table 1 and to make spatial predictions of these variables. Models were usually constructed by backwards stepwise selection, with significance tests for variable removal that varied with model family. Forward stepwise selection was used for species abundance models, because the models generally ended up having few variables included, so that starting with few variables greatly reduced run time for calculating models. Variables were allowed to have either four degrees of freedom, or one degree of freedom.

Separate models were constructed for DOC and non-DOC land, for all variables modelled, except weeds. Spatial predictions were made separately from the two models, and combined at the end for a composite prediction. Finally, predictions were masked to avoid predicting outside the range of the predictor variables in the data.

TABLE 2. SPATIAL PREDICTORS DEVELOPED FOR THIS STUDY.

Not all spatial predictors were used in GRASP because of elimination before modeling for various reasons (see text). Climate and substrate variables are more fully explained in Leathwick et al. (in press). Fragmentation variables are discussed in the text.

NAME OF VARIABLE	ABBREVIATION	DEFINITION	UNITS	CATEGORY	DOC LAND CONDITION	NON-DOC LAND CONDITION	WEED
Mean annual temperature	MAT	Mean annual temperature	C	Climate	Y	Y	Y
Temperature seasonality	TSEAS	The relative deviation of the minimum temperature relative to MAT	C	Climate	Y	Y	Y
Mean annual solar radiation	MAS	Mean annual solar radiation	$MJ/m^2/d$	Climate	Y	Y	Y
Solar radiation seasonality	SSEAS	The relative deviation of the winter solar radiation relative to MAT	$MJ/m^2/d$	Climate			
Rainfall to potential evapotranspiration	R2PET	The ratio of annual rainfall to the annual potential evapo-transpiration	ratio	Climate	Y	Y	Y
Vapour pressure deficit	VPD	The annual vapour pressure deficit	kPa	Climate			
Acid soluble phosphorus	ACIDP	Estimated classes of acid soluble phosphorus	class	Substrate			
Substrate age	AGE	Estimated age class of substrate	class	Substrate			
Soil calcium	CALCIUM	Estimated class of soil calcium	class	Substrate			
Induration	INDURATION	Estimated class of substrate induration	class	Substrate			
Particle size	PARTSIZE	Estimate class of soil particle size	class	Substrate			
Slope	SLOPE	Slope estimated from DEM	degrees	Landform	Y	Y	Y
Elevation	DEM100M	Elevation above seal level	m	Landform			
DOC land (estate)	DOCESTSI	Inside or outside of DOC-administered land	n. a.	Cadastral	Y[1]	Y[1]	Y
Land cover	LCDBSI	Land cover categories from the LCDB	n. a.	Land Cover	Y	Y	Y
Disturbance	DISTURB	Estimated levels of human vegetation removal assigned to LCDB classes	%	Human disturbance			

Disturbance in 300 m radius	DISTURB300	Mean DISTURB in 300 m radius around each pixel	%	Neighbourhood continuous fragmentation	Y	Y	Y
Disturbance in 1000 m radius	DISTURB1000	Mean DISTURB in 1000 m radius around each pixel	%	Neighbourhood continuous fragmentation			
Proportion of DOC land (estate) in 300 m radius	DOC300	The proportion of pixels within 300 m that are in DOC land	%	Neighbourhood continuous fragmentation	Y	Y	Y
Proportion of DOC land in 1000 m radius	DOC1000	The proportion of pixels within 1000 m that are in DOC land	%	Neighbourhood continuous fragmentation	Y	Y	Y
Distance to non-DOC land	LD2NOTDOC	The natural log of the distance of each pixel to the pixel outside DOC [2]	m	Continuous fragmentation	Y	Y	Y
Distance to DOC land	LD2DOC	The natural log of the distance of each pixel to the pixel inside DOC [3]	m	Continuous fragmentation		Y	Y
Size of DOC patch	DOCSIZE	The natural log of the area of the contiguous DOC patch that contains the pixel	ha	Patch-based fragmentation	Y		
Perimeter to area ratio of DOC patch	LOGPTOA	The natural log of the perimeter to area ratio of the contiguous DOC patch that contains the pixel	ratio	Patch-based fragmentation			

[1] Used as basis for separating region into DOC and non-DOC areas.

[2] For pixels that were outside DOC land (estate), the distance was set to 25 m to solve problems associated with log transformations.

[3] For pixels that were inside DOC land, the distance was set to 25 m.

For each modeled variable, the following analyses were produced:

1. Graphs of the modelled variable plotted against each candidate spatial predictor. For 0/1 variables (i.e. presence/absence variables), histograms of the proportion of 1s were graphed against each predictor variable. For continuous response variables, scatterplots were made of the modelled variable against each predictor variable.
2. The final GAM model, with the curve of the partial contribution of each predictor variable to the overall model.
3. The model validation and cross-validation results, showing the plots of the observed versus the predicted values for each. For continuous variables, the correlation between predicted and observed was used to assess the model. For binomial models, the Receiver Operating Characteristic (ROC) statistic (Fielding & Bell 1997) was used.
4. The relative contributions of each spatial predictor variable to the model. These were done both as alone contributions and as drop contributions to the model. The alone contribution of a spatial predictor variable shows the deviance explained by a model with only that variable in the model. The drop contribution of a spatial predictor variable is the difference in explained deviance between the full final model and a model with the variable excluded. If the variable in question is not in the final model, the drop contribution is defined as zero.

3.6.1 Sampling densities and plot weights

Models of plot distributions were used to predict the sampling density for the entire study region. The surface of predicted sampling density was treated as the best *post hoc* estimate of the inclusion probability for each sample point. In all subsequent analyses, the estimated inclusion probabilities were used to weight observations. The weights were defined as the inverse of the sampling probabilities. To keep the effective sample size equal to the number of plots, the weights were standardised by the mean weight to achieve a mean weight over all samples of 1.0. The overall effect of this weighting was to up-weight the samples that were in sparsely sampled regions, and to down-weight the samples in heavily sampled regions.

Sampling densities were modelled in relation to spatial predictor variables using binomial models. Separate models were made for inside and outside DOC-administered land (estate). A 1% systematic sample of the pixels of the study area was chosen as absences and assigned a weight of 100. All pixels containing a sample were defined as presences and assigned a weight of 1.0. All weights were divided by the mean weight over the presences and absences to achieve an overall weighting of 1.0. A chi-square test was used to test variable significance, and all variables with $p < 0.01$ significance level were removed from the model.

3.6.2 Vegetation condition index models

Vegetation condition indices were modelled in relation to spatial predictor variables using the models shown in Table 1.

3.6.3 Weed species models

Weed distributions were modelled using a composite model, using methods similar to Welsh et al. (1996) and Pearce & Ferrier (2001). The first part of the composite model described the probability of presence or absence of the species. The second part of the model used only the points in which the species was present, and modelled the cover of the species, given that it was present. The predictions from each of these models were multiplied together to give a composite model of weed abundance. In some cases, the abundance model failed, in which case the weed distribution model was simply a presence-absence model.

The sample size for the presence-absence models was all the plots, while the sample sizes of the conditional abundance models was only the number of plots in which the species occurred and the cover was recorded (Tables 1 and 4).

3.6.4 Prediction masks and combining predictions

All predictions were masked to avoid predicting outside the range of the data. Separate masks were created for DOC and non-DOC land, and then applied separately. Masks were defined by the range of each spatial predictor spanned by the data. Pixels in the prediction grid that fell outside the range of the data on any of these axes were masked out. While it might be possible to predict slightly outside the ranges of the observed data, this was not done here, because sample densities tend to decrease towards the edge of the distributions, increasing model uncertainty at the edges of the range and making further extrapolation unwise.

4. Results and discussion

4.1 GENERAL

The study area was 5 723 195 hectares, with 1 974 446 ha (34%) managed by DOC for conservation purposes (called DOC land in this report). The collection of past surveys was dominated by surveys of indigenous vegetation, especially forest, resulting in a strong bias of the data towards areas where the vegetation condition is good (Fig. 4). Thus, while existing surveys were perfectly valid for their intended purpose, special techniques must be applied to use these data for other purposes. Similar problems have been faced by the carbon monitoring work (Hall et al. 2001) in estimating carbon storage, with a different set of solutions used. The total number of plots was 16 672. This consisted of 1326 grassland transects, 302 recce plots from Landcare Research 2000 surveys, 328 recce plots from the surveys associated with this project, and 14 716 recce plots from past surveys. Of the past recce surveys, 8721 plots (59%) had cover scores. This leads to a total of 9351 recce plots with cover used to model variables requiring cover information. Of the total number of plots, 14 272 (86%) were on lands managed by DOC, and 2400 were on non-DOC land.

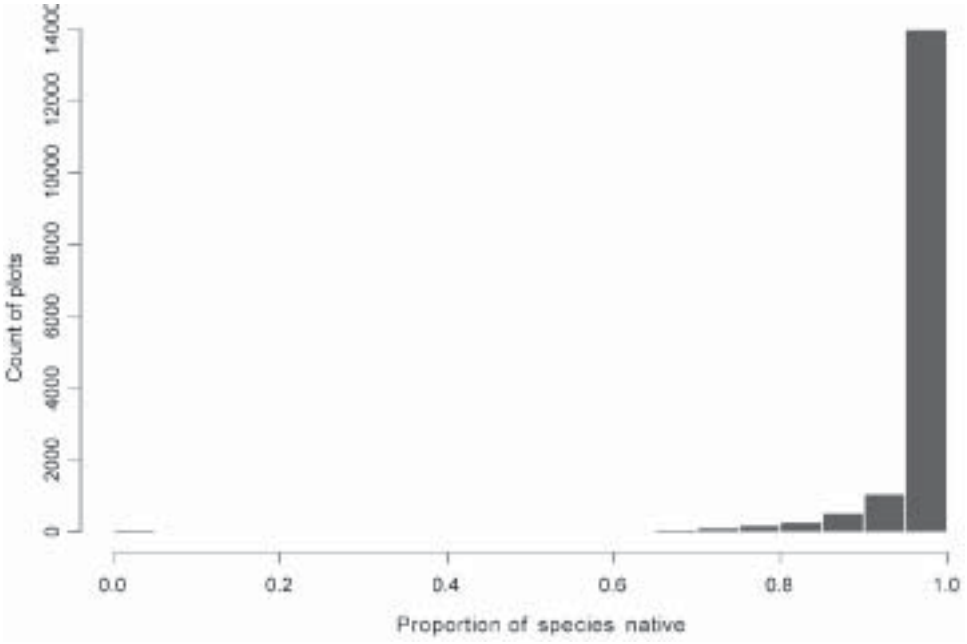

Figure 4. The distribution of a condition index across plots. The histogram of the proportion of species native is dominated by the large number of plots in indigenous vegetation. This highlights the difficulty of using data from past surveys for current applications.

The dates of the vegetation data span more than 40 years, from the 1960s to the present (Fig. 5). The majority of the large numbers of recce plots in indigenous forest and grassland transect data were gathered in the 1970s and early 1980s. The surveys in 2000 and 2001 spanned a much wider range of vegetation, e.g. fell-fields, pasture, plantation forestry and urban areas. The 2001 surveys were specifically designed to sample some of the under-sampled land uses and areas.

4.2 FIELD SAMPLING

A number of improvements could be made if more time was available for planning and coordinating vegetation sampling. A number of difficulties were encountered. These included sampling not starting until well into a very dry summer, so that many annual plants had withered before sampling took place.

While attempts were made to have consistent sampling methods, there were still variations resulting from differences in interpretation between individuals. Consistency of plot methods would be improved if all field teams were collected at the beginning of the season and surveyed plots together (L. Burrows, pers. comm.).

For various reasons, it was not always possible to get to all the targeted sampling locations. Generally, this was because the location was dangerous or particularly inaccessible. This inability to sample all locations cannot be avoided; its effect on the overall results is unknown, but probably small.

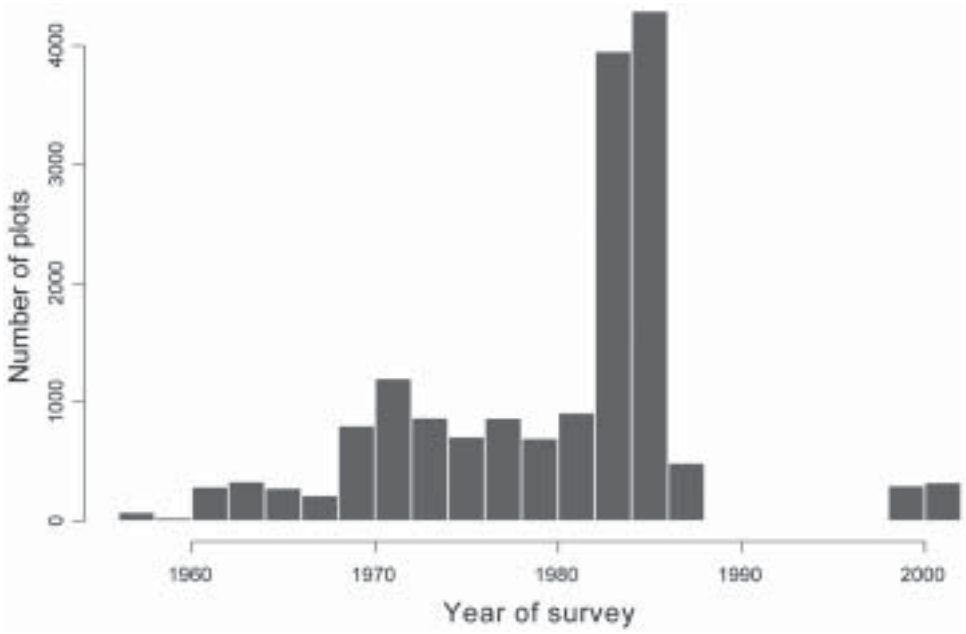

Figure 5. The distribution of year of plots across all samples.

4.3 MODELED PLOT PROBABILITIES

The overall sampling probability can be calculated nonspatially by dividing the number of plots by the number of hectares. This gives plots per hectare, or the average probability that any one hectare is sampled. Since the probabilities are quite small, it is often easier to discuss the inverse of the probability, or odds of sampling (for example: 1 hectare sampled out of 500 hectares). Technically, 'odds' is usually defined as p/(1–p), so the term is used loosely here for lack of a better word.

For the DOC land, the total area was 1 974 446 hectares, and the total number of plots in this area was 14 272. This leads to an overall average probability of sampling of 14 272/1 974 446 = 0.0072, or one plot sampled in 138 hectares (1:138 odds). The mean predicted odds of sampling was close to this, at one plot in 132 hectares (1:132 odds), with significant spatial variation, from 1:5900 to 1:31 odds. The model and contributions are shown in Figs. A1 and A2.

The total area of non-DOC land was 3 748 749 hectares, with a total 2400 plots. This leads to an overall average probability of sampling of 2400/374 879 = 0.000640, or 1:1563 odds. The mean over the predicted grid of the probability of sampling was close to this at 0.000698 or 1:1433 odds. The predicted odds of sampling showed a wide spatial variation from 1:50 000 to 1:45. The model and contributions are shown in Figs. A3 and A4.

The combined prediction of sampling probability is shown in Fig. 6. Here the probability of sampling is given as odds of sampling. This surface is used to predict the sampling probability for each sample point. The inverse of these predicted probabilities is used to weight the observations in all subsequent analyses.

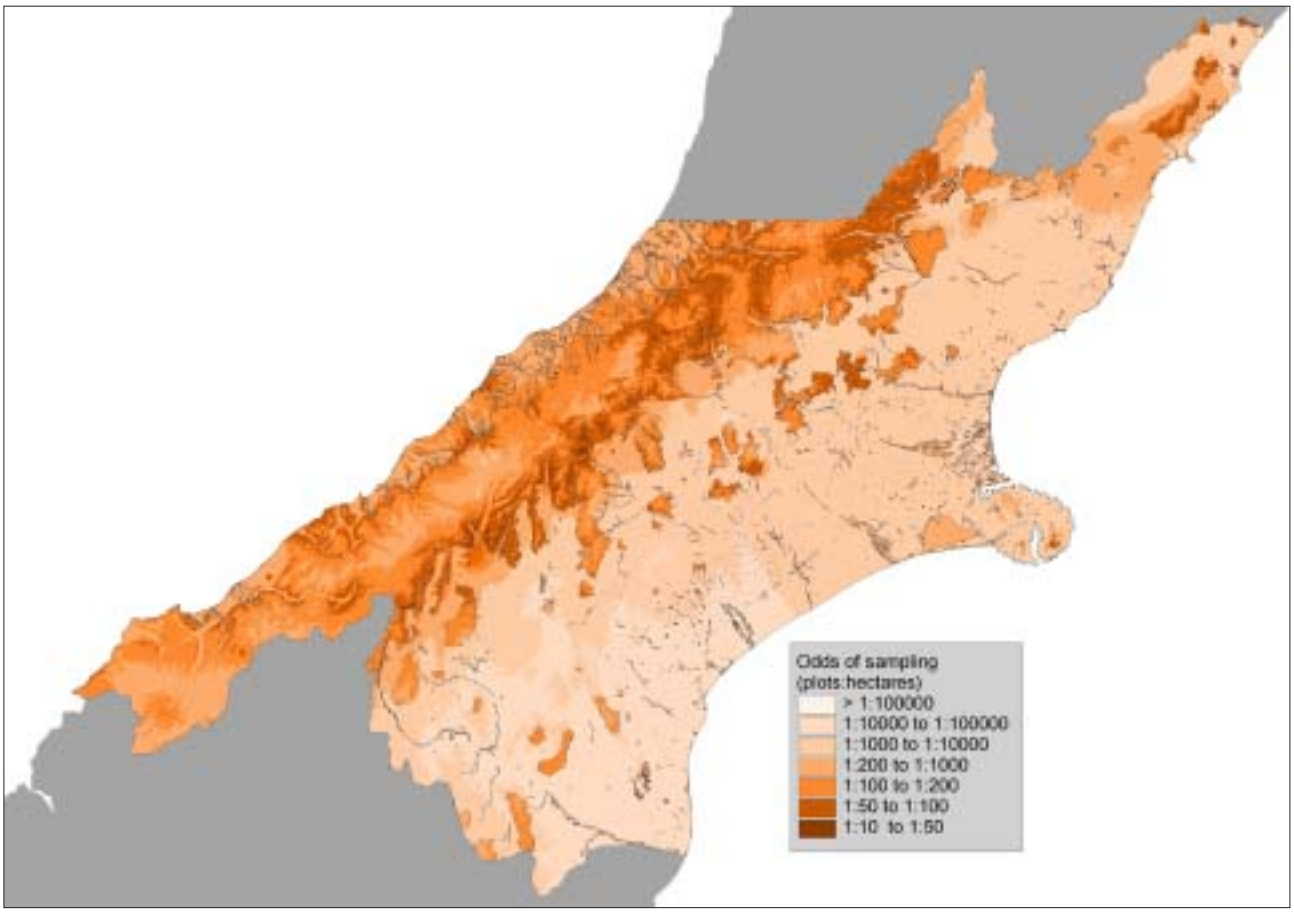

Figure 6. The modelled density of vegetation plots across the study area. The map results from modelling the distribution of sample points seen in Fig. 3. Separate models of plot density were calculated for inside and outside DOC-administered land (boundary shown in black). The odds of sampling are defined as one divided by the probability that a hectare contains a plot. The odds of sampling are the number of hectares that would be needed on average to contain one plot. When the original sample plots are overlain onto this map, the odds of sampling are particularly useful because they predict the number of hectares that each plot represents in the landscape. Plots in sparsely sampled portions of the landscape represent more hectares than plots in heavily sampled areas. These predicted odds are used to weight samples in the analyses.

4.4 DATA UNCERTAINTIES

While the modelled plot probabilities give the overall uncertainties due to sampling intensity, there are a number of other uncertainties that affect the data. Plot data have various positional, recording and species identification errors, as well as random and systematic variation between observers in estimating cover. While attempts have been made to reduce and understand these errors, they remain poorly quantified.

The spatial information used as predictors (Table 2) also has associated uncertainties. However, this is easier to deal with in these methods, since the uncertainty of the spatial predictor simply degrades the relationship between the predictor and the variable of interest. This uncertainty of the relationship is indicated by the statistics that describe the strength of the regression, such as the D^2 statistic used for the condition indices.

For instance, it is clear that the LCDB has many errors. However, the fact that the LCDB is an imperfect prediction of the actual land cover does not mean that it cannot be used, since this will be reflected in the strength of the ANOVA embedded in the GAM regression model. The strength of the model will depend on the strength of the relationship between the modelled variable and true land cover, degraded by the inaccuracy of the LCBD in predicting the actual land cover.

4.5 FRAGMENTATION

A wide range of indices and analyses exist for analysing fragmentation and the spatial pattern of landscapes (e.g., Gustafson 1998). Patch-based indices of fragmentation focus on 'patches' of habitat, surrounded by unsuitable habitat. Patch-based fragmentation indices generally use polygons to depict the 'good' patches, and consider various characteristics of the polygons, such as their size, shape, and distance to other patches. This focus on patches and polygons produces both operational and conceptual difficulties.

The conceptual difficulties of patch-based statistics come from the artificial depiction of the landscape into 'good' patches surrounded by 'bad' matrix. This strict 'good' and 'bad' division of the landscape has difficulties when applied to a wide variety of environments and land uses, as in this project, since the aim is to understand native vegetation condition across all land uses and vegetation types, not just across a single type, such as indigenous forest.

The operational difficulties of patch-based fragmentation statistics are also considerable:

1. Patch statistics are extremely dependent on mapping resolution and accuracy of the data used to define the patches. For example:

   a) If small habitat fragments are ignored when drawing the polygons that define the patches, the result will be a different patch size distribution than there would have been had the small fragments also been mapped.

   b) If a road going through the middle of a large patch is mapped, the patch size is suddenly half of what it would have been without the road.

   c) If an isthmus of forest (for example, in Lees Valley) is mapped as connected to a large forested area, such as the forests of the central Southern Alps, its patch size will be very large. If the area is mapped somewhat differently, and is not connected, then its patch size will be very small.

2. The isthmus effect of 1c above exemplifies a fundamental problem with the patch-based approach. Does it really matter if the patch in Lees Valley is connected so that you could walk to Fiordland without leaving forest? While this might be important for a few processes, for most matters of ecological interest, more local effects will prevail. For example, that the isthmus of forest is primarily surrounded by pasture or gorse, and all the disturbances associated with these land covers, is much more important than whether it is mapped as being connected to other places, such as the Routeburn Track.

3. Patches are not just 'good' or 'bad'. Native forest surrounded by scrub might be in better condition than native forest surrounded by pasture. Patch definition can become confusing and arbitrary, if there are many varying habitat types.

These problems greatly reduce the usefulness of patch-based fragmentation indices. Because of the wide range of land covers spanned in this study, it was not reasonable to calculate patch-based fragmentation indices for each different land cover type, so these were calculated only for polygons of the DOC land. Both DOCSIZE and LOGPTOA (see Table 2 for definitions) have very skewed distributions across the data, because some very large polygons extend through the large tracts of forest in the Southern Alps.

Because of these problems with patch-based measures of fragmentation, we developed continuous measures of fragmentation that deal better with the continuous nature of habitat variation and are less sensitive to details of mapping. These measures include:

1. Distance to an edge. This measure provides a continuous measure of edge effect. The edge could be a certain type of land cover or land tenure. This measure reduces problems 1 and 2 above significantly, but still suffers from Problem 3. Since this study was concerned with all land-cover types, land cover was not used to define edges. Doing so would have required calculating surfaces of distances to each land type, such as distance to forest, distance to pasture, etc. Instead, a surface of the distance to the edge of the DOC land and a surface of the distance to the non-DOC land were calculated. Variables of this type used in this report include LD2DOC and LD2NOTDOC (Table 2)
2. Proportion of the habitat type within a neighborhood. This type of measure gives the proportion of the pixels within a neighborhood (often a circle of a given radius) that are of a particular type of land cover or land tenure. The overall result is that the map cover or tenure type is essentially fuzzed out. As in the distance measures, this measure does much to reduce problems 1 and 2 above; it is much less sensitive to the details of mapping, and also treats an isthmus essentially the same as an island. However, this measure also still suffers from problem 3, since it focuses on discrete patch types. For this reason, it was not calculated for different land covers, but solely for the proportion of pixels in the neighborhood in DOC land. Variables of this type used in this report include DOC300 and DOC1000 (Table 2).
3. Disturbance intensity within a neighborhood. This is similar to problem 2 above, but also avoids discrete patch types and therefore addresses problem 3 as well. Here, each cover class is assigned a disturbance level, and a neighborhood mean is then calculated around each pixel. Thus, this measure captures whether a pixel of forest is surrounded mostly by low disturbance types like forest, or higher disturbance types like pasture, and can easily deal with mixes of many different land covers. Variables used in this report include DISTURB300 and DISTURB1000 (Table 2).

All these measures give a gridded characterisation of fragmentation that varies continuously across the landscape. They avoid some of the problems of the patch-based statistics, and have more tractable statistical distributions.

Interpretation of the fragmentation variables and their contributions to the models is difficult. Certainly, some of the effect of the fragmentation variables is to take into account positional errors in both plot locations and boundaries, for instance in legal boundaries of parcels of DOC land. This is important because it was clear that a plot designed to be in DOC land sometimes fell just outside it. In this case, the fragmentation variables are helping by dealing with some of the inaccuracies of the data, not because of ecological processes generally associated with fragmentation.

TABLE 3. STATISTICS FOR THE MODELS OF VEGETATION CONDITION. SEPARATE MODELS WERE CALCULATED FOR AREAS INSIDE AND OUTSIDE DOC-ADMINISTERED LAND. FOR EACH MODEL, THE DEVIANCE EXPLAINED (D^2) AND APPROXIMATE DEGREES OF FREEDOM (d.f.) ARE SHOWN.

CONDITION INDEX	OUTSIDE DOC LAND		INSIDE DOC LAND	
	D^2	d.f.	D^2	d.f.
Native species richness	0.608*	41	0.510*	50
Exotic species richness	0.653*	32	0.496*	49
Proportion of species native	0.830	39	0.676	46
Cover	0.608	22	0.395	50
Proportion cover native	0.806	28	0.801	40
Biomass index[1]	0.641	35	0.500	54
Proportion biomass native	0.813	32	0.830	42

* Warning: model unstable under cross-validation due to rare, extreme predictions.

4.6 CONDITION INDICES

Summary statistics for the models of all condition indices are given in Table 3. Land cover was the most important predictor of vegetation condition, with environmental variables also important, and fragmentation variables having a lesser (but statistically significant) importance (Fig. 7). The models were generally successful in predicting the indices, with only the models for richness variables unstable under cross-validation due to rare, extreme predictions. Instability under cross-validation suggests that the model might have difficulties if tested on new data, and the predictions should be treated with caution. For five out of seven of the indices, the deviance explained was greater for non-DOC land. This is presumably because the better overall condition on DOC land leads to less overall variation in condition. The deviance explained is also sensitive to sample design, and the biased distribution of plots has an influence on these statistics that is difficult to quantify. The spatial prediction for the proportion of species native is shown in Fig. A5, with the models and contributions shown in Figs A6 to A9. The spatial prediction for native species richness, combined for DOC and non-DOC land and masked to avoid predicting outside the data range. is shown in Fig. A10. The models and contributions are shown in Figs A11 to A14.

The construction of separate models for DOC and non-DOC land required considerable extra effort, but had a number of advantages. First, since conservation lands are specifically managed and protected for conservation, it is reasonable to expect the relationships between vegetation condition and environmental characteristics to be different inside and outside these lands. Second, this project was developed specifically for DOC, and the relationships within DOC land are of particular interest. Models of the entire landscape would mix the two effects, and would be more difficult to interpret. Finally, the model could only be improved by splitting, as this would provide extra flexibility for the model to fit the data.

While other measures of condition may be used (e.g. see Allan 2000), these results exemplify the use of GRASP to predict biodiversity condition for use in systematic conservation management. Condition is a major component of site

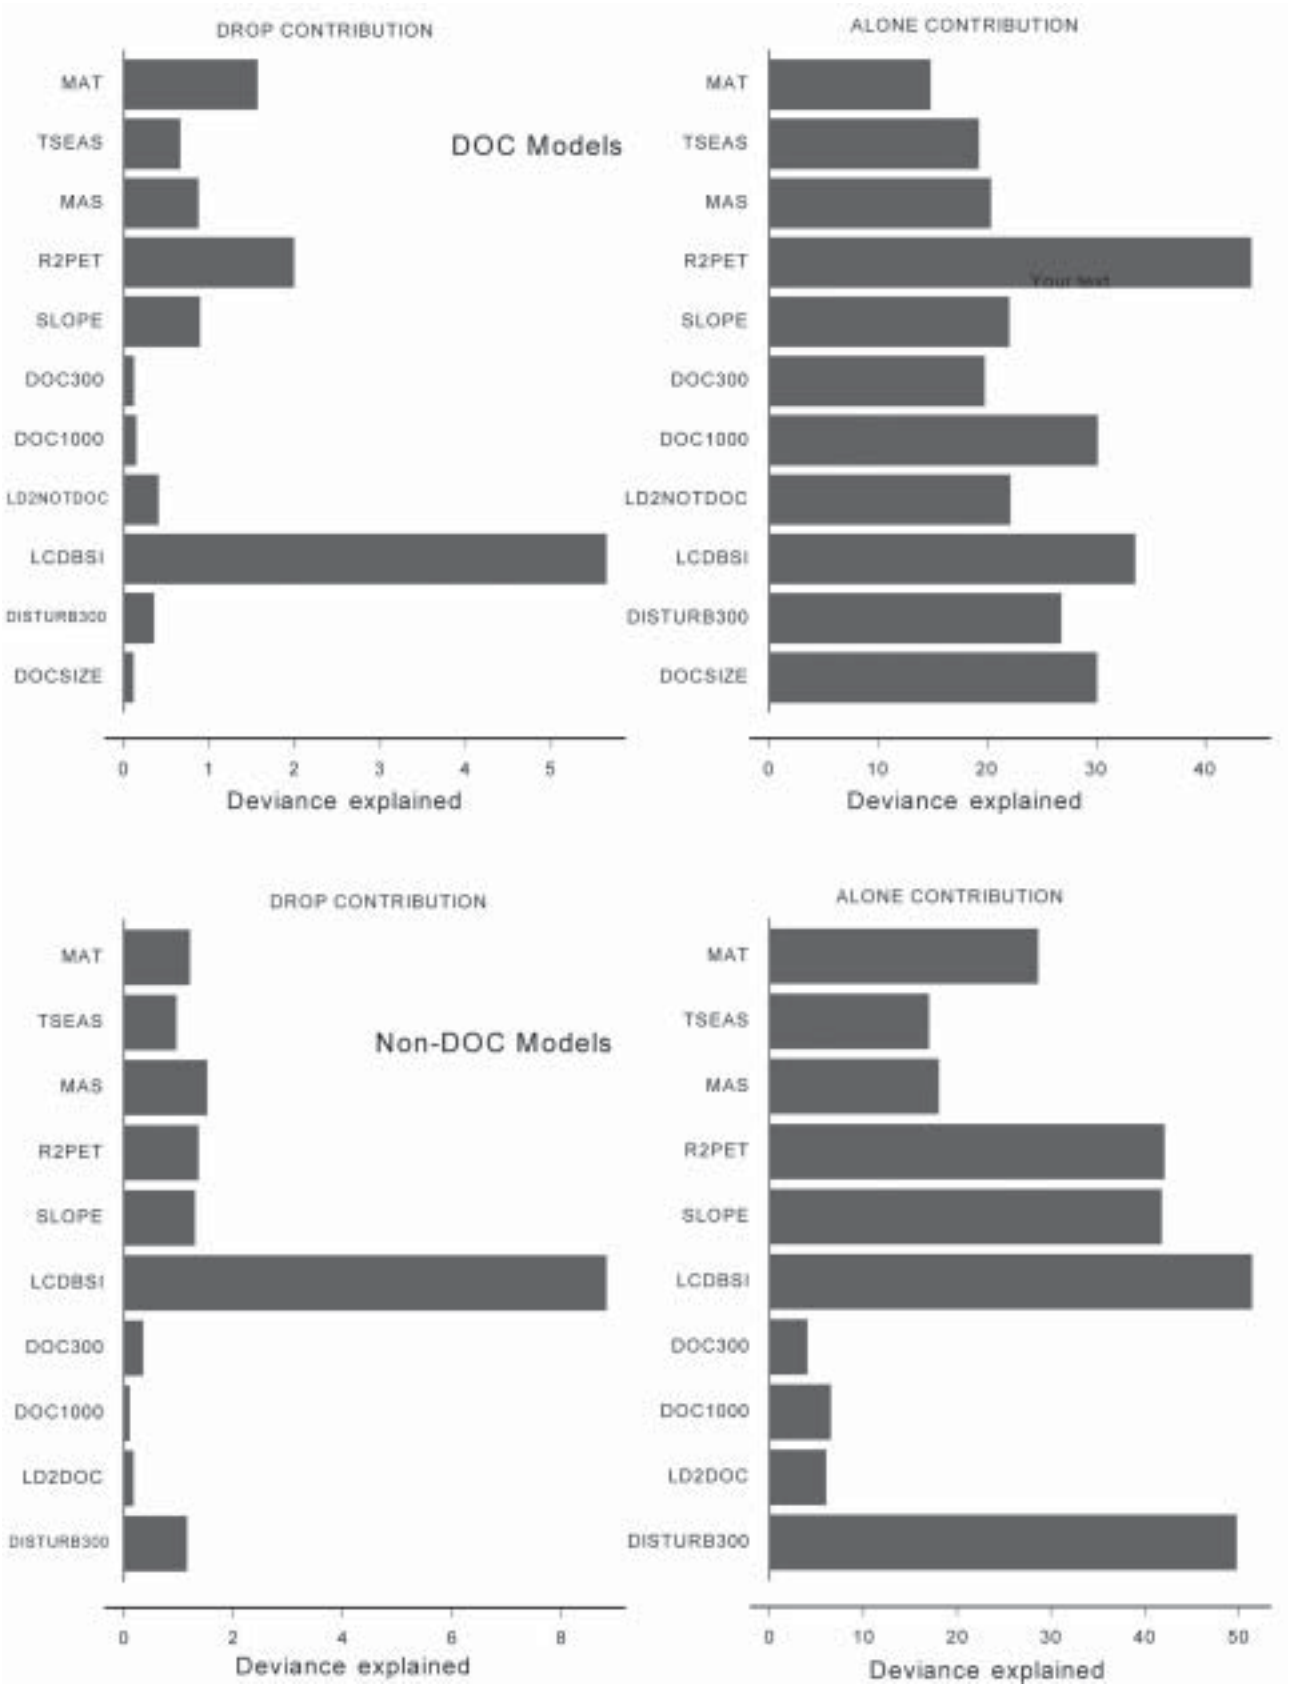

Figure 7. The strength of spatial variables in predicting vegetation condition. This figure shows the deviance explained by each spatial predictor variable, averaged over all condition indices. DOC and non-DOC land models are shown separately. See Section 3.6 for an explanation of 'alone' and 'drop' contributions. The contributions to biomass and cover were excluded from these analyses, since they do not relate directly to indigenous condition.

value, and the one that is most affected by conservation management. Quantifying the difference made by conservation action requires spatially explicit predictions of condition for entire regions. Since condition will never be measured for every part of the region, the only way to achieve this is by some sort of spatial prediction technique, such as demonstrated here.

4.7 WEED DISTRIBUTIONS

The weed models are summarised in Table 4. Of the 37 weed species modelled, 21 (57%) had successful abundance models and were thus modelled with composite models. Some of these models should be treated with caution (as noted in comment column of Table 4), either because the cross-validations were driven by large values, or because of significant components of unidentified congeners in data. Of the remaining 16, the abundance models failed for the following reasons: model problems indicated by cross-validation (10); no significant variables (4); no plots with cover data (1); problems with model (1). Most of these problems can be traced to low numbers of plots with cover data, leading to small sample sizes in the abundance models. Predicted distributions of all weed species are shown in Fig. 8. A more detailed plot of the composite spatial prediction for gorse, *Ulex europaeus*, is shown in Fig. A15, with the models and contributions shown in Figs. A16 to A19.

A number of genera, such as *Senecio* and *Trifolium*, were not modelled, because of large numbers of records of unidentified species. The genus *Hieracium* also had a significant component of unidentified species, but three species of *Hieracium* were modelled anyway, because of interest in them. Other genera also had unidentified species, but were usually 1–10% of the overall records.

While, overall, the weed models performed well and captured the pattern of the species, there were some details that were not perfect. For instance, detailed examination of the spatial predictions of presence-absence for all *Pinus* spp. (PINUSALL) shows that it is not always higher than the predictions for *Pinus radiata* alone, which it should be, since PINRAD is a component of PINUSALL.

There is also reason to question the fine scale details for all *Salix* spp., although examination of the predictions against the data suggests that the model is fitting the data quite well. While it is difficult without additional observations to fully understand the problems, some of the patterns for willows along rivers are questionable. There are also predictions for willows in some of the forestry areas on the Canterbury Plains; however, there are sample points with LCDB classes of plantation forestry that include willows.

The composite models used to model weed distributions were generally very successful, and had both statistical and ecological appeal. While skewed distributions are commonplace in species distribution data, the fact that the data used in the model spanned very different environments and land uses meant that the problem with skewed distributions was particularly pronounced. The distribution of abundances was very highly skewed, with many plots having zero cover, and a few plots with medium or high values. The composite models dealt well with this situation by first modelling the position

TABLE 4. WEED SPECIES OCCURRENCES AND MODEL TYPES. THE SIX-LETTER CODE USED IN NAMING GRIDS IS SHOWN, ALONG WITH THE NUMBER OF PLOTS IN WHICH THE SPECIES WAS PRESENT, AND THE NUMBER FOR WHICH IT WAS PRESENT AND COVER WAS RECORDED. THE MODEL TYPES ARE EITHER A COMPOSITE MODEL OF A PRESENCE-ABSENCE AND ABUNDANCE MODEL OR SIMPLY A PRESENCE-ABSENCE (PA) MODEL. IF THE ABUNDANCE MODEL WAS NOT USED, THE REASONS ARE GIVEN UNDER COMMENTS. THE MOST COMMON REASON IS THAT THE CORRELATION BETWEEN OBSERVED AND PREDICTED VALUES UNDER CROSSVALIDATION (CVAL) WAS TOO LOW FOR THE ABUNDANCE MODEL.

SPECIES	SPECIES CODE	PRESENCE	COVER	MODEL TYPE	COMMENTS
Agrostis capillaris	AGRCAP	365	547	Composite	
Anthoxanthum odoratum	ANTODO	522	800	Composite	
Cytisus scoparius	CYTSCO	63	73	PA	cval < 0.1
Festuca rubra	FESRUB	124	235	Composite	Caution, cval driven by large values
Hieracium lepidulum	HIELEP	11	12	PA	cval unsuccessful, 174 records of unidentified *Hieracium* sp.
Hieracium pilosella	HIEPIL	206	263	Composite	Caution, high preds for zeros
Hieracium praealtum	HIEPRA	262	352	Composite	Caution, extreme prediction in cval
Holcus lanatus	HOLLAN	419	727	Composite	
Lupinus arboreus	LUPARB	17	17	Composite	
Pinus radiata	PINRAD	50	51	PA	cval = 0.12 only MAS in model
Pinus species combined	PINUSALL	61	66	Composite	
Rosa rubiginosa	ROSRUB	35	45	Composite	Caution, cval driven by large values
Rubus fruticosus	RUBFRU	47	50	PA	cval < 0.1
Salix species combined	SALIXALL	14	14	PA	cval < 0.1
Ulex europaeus	ULEEUR	117	122	Composite	
Additional weeds					
Agrostis stolonifera	AGRSTO	52	53	PA	cval = 0.1
Aira caryophyllea	AIRCAR	35	53	Composite	
Bromus hordeaceus	BROHOR	48	48	PA	No sign. vars.* for abundance model
Cerastium fontanum	CERFON	77	232	Composite	
Cirsium arvense	CIRARV	75	89	PA	No sign. vars. for abundance model
Cirsium vulgare	CIRVUL	148	218	Composite	
Crepis capillaris	CRECAP	193	337	Composite	
Dactylis glomerata	DACGLO	165	286	Composite	
Digitalis purpurea	DIGPUR	42	65	Composite	
Hypochoeris radicata	HYPRAD	596	1072	Composite	
Juncus articulatus	JUNART	52	52	Composite	
Juncus effusus	JUNEFF	49	49	PA	cval driven by large values
Lapsana communis	LAPCOM	268	269	Composite	
Linum catharticum	LINCAT	78	138	PA	cval < 0.1
Lolium perenne	LOLPER	101	107	Composite	
Lotus pedunculatus	LOTPED	103	104	PA	cval < 0.1
Mycelis muralis	MYCMUR	266	810	Composite	
Poa annua	POAANN	35	110	PA	No sign. vars. for abundance model
Ranunculus repens	RANREP	55	68	PA	Contradictory fragmentation vars in ab model
Rumex acetosella	RUMACE	265	575	PA	No sign. vars. for ab model, 54 records of unidentified *Sonchus* sp.
Taraxacum officinale	TAROFF	86	170	PA	cval < 0.1
Verbascum thapsus	VERTHA	51	56	PA	cval < 0.1

*sign. vars. = significant variables

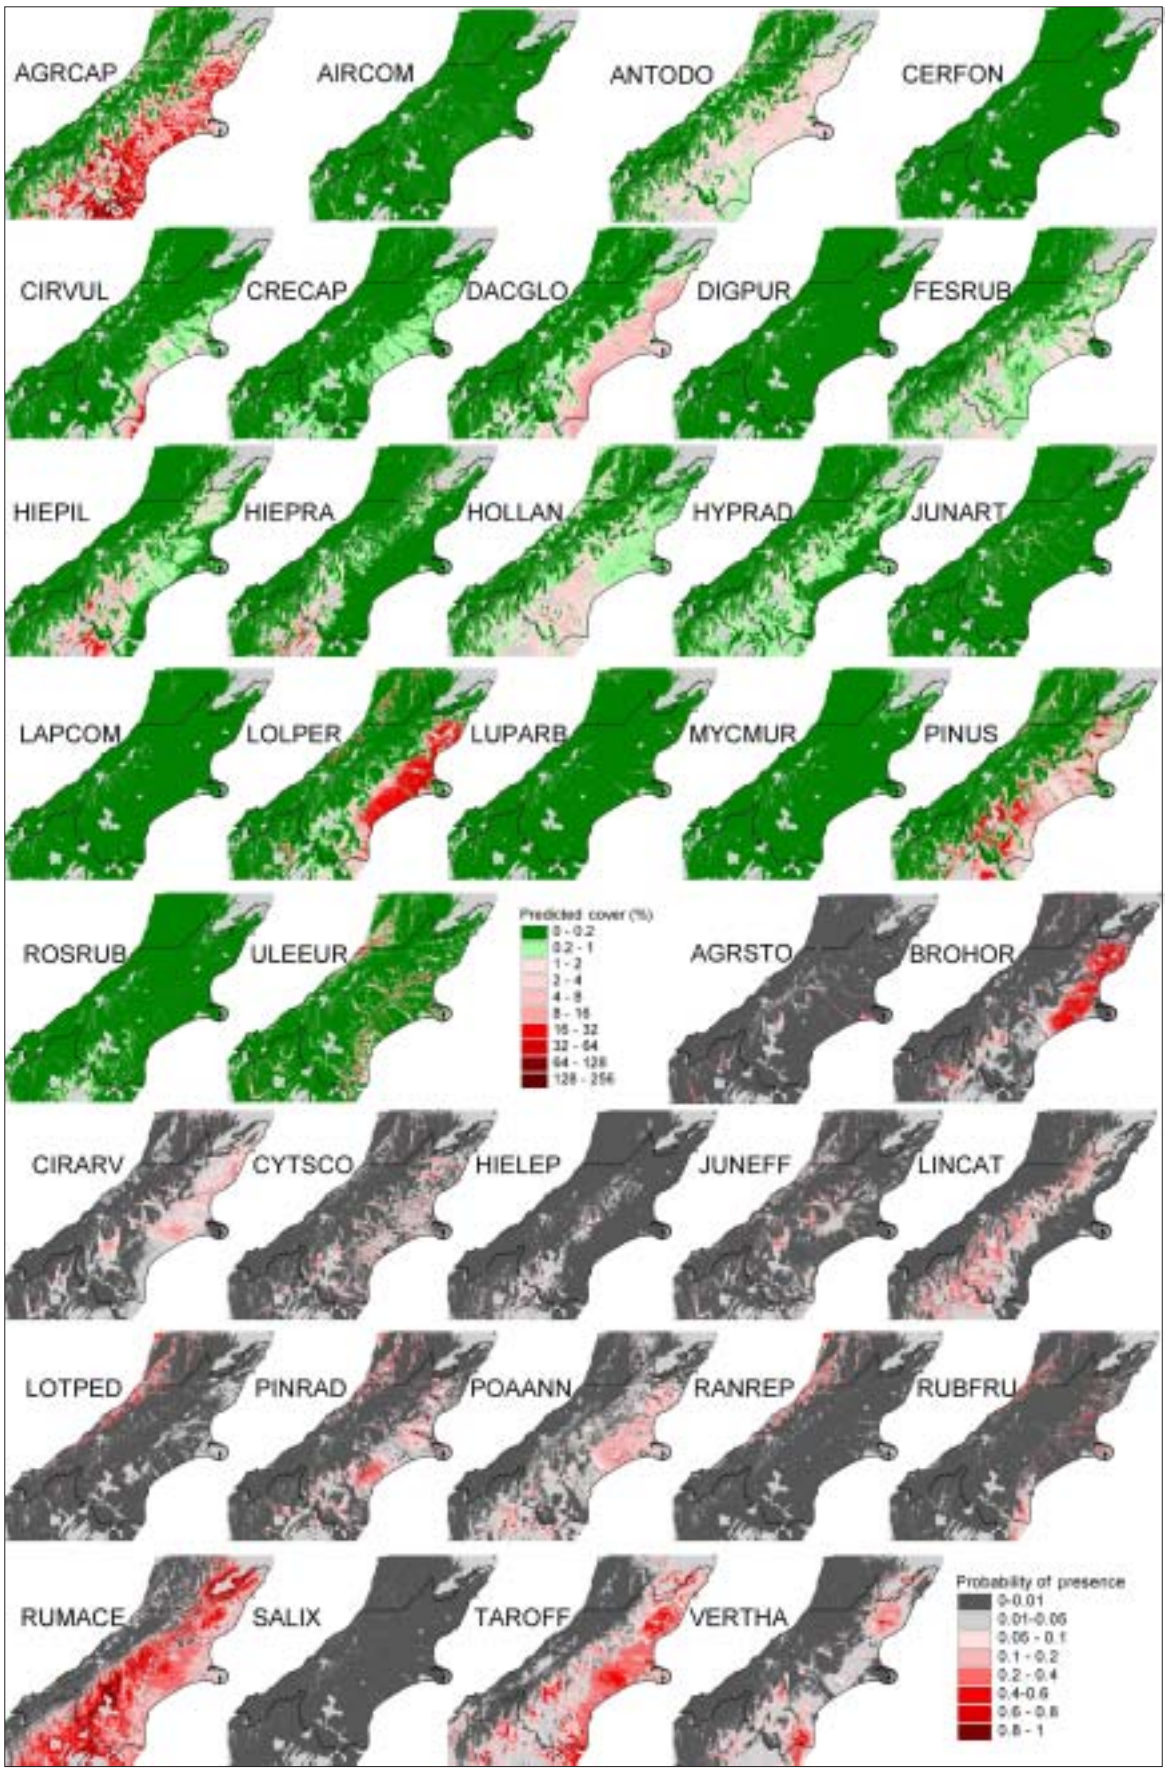

Figure 8. Predicted distributions of 37 weed species. The green and red colour scheme is used for composite models, showing a predicted percentage cover. Since the cover is summed over a number of tiers, it can go over 100%. The grey and red colour scheme is used for presence-absence models, and shows the probability that a species is present in a plot.

of the presences, then modelling the abundance within those presences. This model is also a good reflection of the most likely ecological reason for the skewed distribution of species abundances—that quite different processes affect habitat suitability (i.e. whether the species is present) than those that determine its abundance within suitable habitat. Finally, the composite model used the data efficiently, because all plots were used to model presence-absence, while often less than 100 points were used to model abundance.

Overall, these predictions of weed distribution and abundance provide spatially explicit estimates of conservation pressures for large areas. These are used explicitly in MCA to assess pressures on and condition of sites. They are also the beginning point for predicting the spread of these weeds over time, under different management scenarios. The difference between the distributions of uncontrolled weeds, and their effects on native biota, and weed distributions under different control scenarios is a measure of the difference made by conservation action.

4.8 CORRECTION OF SAMPLE BIAS

In this study, the bias in the distribution of sample points was investigated and corrected using two techniques. The modelling of sample density (Sections 3.6 and 4.3) showed a significant variation in sampling density across the study region (Fig. 6), with denser sampling in areas of high indigenous vegetation condition. If a simple mean of the plots was used to estimate the mean vegetation condition for the region, then it would indicate that the condition of the study region was much better than it is in reality. The first correction of this bias was affected by weighting each plot by the inverse of its estimated inclusion probability (see Section 3.6). This gave greater weight to plots in sparsely sampled areas. The second correction is inherent in the regression methods used by GRASP. The process of producing models and predicting them onto grids means that plots are only used to predict onto the parts of the landscape that they represent. Because of the combination of these two corrections, the means of the predicted grid are much more accurate predictions of condition for the region than are the means of the plots.

Overall, the plots had a higher mean vegetation condition and lower mean weed abundance than the predicted grids, illustrating the dangers of inferences based on found data. This is illustrated in Fig. 9, using a single condition index (proportion of species native) and a single weed species (gorse). The plot means are dramatically higher in both of these characteristics for the overall study region and non-DOC land, and slightly higher for DOC land.

4.9 ADVANTAGES AND LIMITATIONS OF GRASP

Systematic conservation management in general, and the MCA process in particular, requires spatially explicit descriptions of the pressures, such as animal and plant pests, affecting conservation and other land, and reporting requires estimates of biodiversity condition, change in pressure and change in condition. GRASP is a leading method for providing these predictions in a rigorous and quantitative manner.

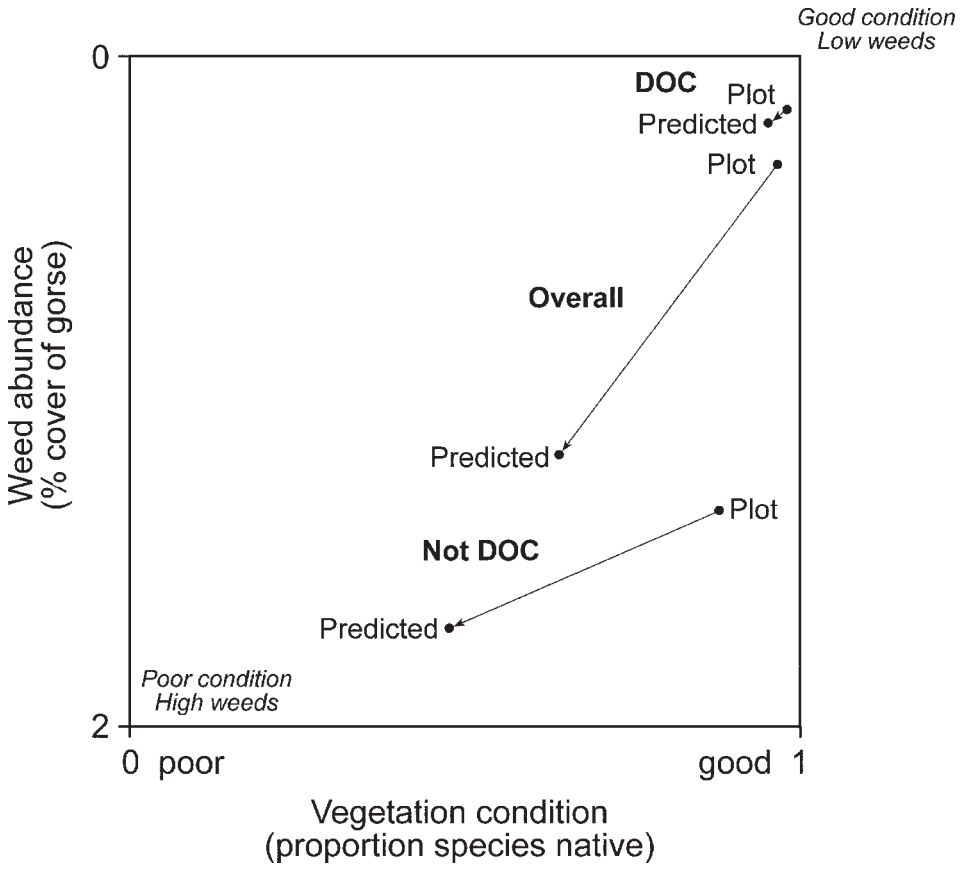

Figure 9. The correction of bias afforded by GRASP. The two axes of condition and weed abundance show areas of higher indigenous value in the upper right, and lower indigenous value in the lower left. A pair of points is shown for a) the overall study area, b) DOC and c) non-DOC portions of the study area. If a simple mean of the plots was used to assess the vegetation condition or weed abundance, then the points labeled 'plot' would be the estimated condition and weed abundance for the area. When the GRASP process is applied, and predicted onto the grid, this corrects for the bias in plot distributions. The points labeled 'predicted' give the mean of the predicted grid for condition and weed abundance. The difference between them (arrows) show the correction in bias provided by GRASP. In all cases, the plot means are biased towards lower weed cover and higher condition, especially in the non-DOC areas, providing erroneously high estimates of indigenous value.

The results of this study demonstrate the power and utility of GRASP (for the MCA process). Using the index proportion of species native (PSppNat) as an example, the process has provided a spatially explicit prediction of PSppNat across the region (Fig. A5), derived from objective plot-based measurements. The models of PSppNat show the different patterns of native vegetation and influences on it for DOC and non-DOC lands. The utility of GRASP in dealing with biased data is evident in Fig. 9. A simple mean of the plots would have given a much higher overall mean for the non- DOC land than is predicted using GRASP.

The biggest limitation of GRASP, and other empirical techniques, is the availability of, and access to, quality data. Much of the cost of implementing these techniques lies in providing the data they require.

5. Recommendations

5.1 INVEST IN RIGOROUS MONITORING SYSTEMS BASED UPON PROBABILITY SAMPLING SCHEMES TO PROVIDE THE INFORMATION REQUIRED BY GRASP TO SUPPORT THE MCA PROCESS

If conservation planning and reporting is to be based upon robust and defensible information, then an investment is required in the underlying information. Scattered ad hoc data collection may be sufficient for one-off or local purposes, but accurate and reliable results at regional and national scales requires coordinated and planned sampling, data collection and analysis. This can only be provided by data that is gathered according to rigorous probability sampling schemes, using well-designed measurement techniques.

5.1.1 Advantages of probability sampling

Sampling needs to be carried out rigorously, according to a defined probability sampling scheme. Examples of probability sampling include simple random sampling, most grid sampling, or stratified random sampling. The biggest advantage of probability sampling schemes is that they allow unbiased and defensible results. While the methods used here to account for the bias of found data have been internationally reviewed, a better solution is to design reliable sampling schemes from scratch.

5.1.2 Grid versus stratified sampling

Grid and stratified sampling are not mutually exclusive. Grid sampling is an excellent method to provide an unbiased and overall coverage of a region. However, with limited budgets, there can be considerable advantage in increasing efficiency by, for example, sampling more intensively in areas of higher interest. A grid design actually maximises the amount of travel required to collect samples, and reducing travel time provides opportunities to gather similar information for less cost.

Spatial information available in geographic information systems presents a powerful opportunity for designing sampling schemes that are rigorous and defensible as well as efficient in sampling rare or important areas and in reducing travel costs. Cawsey et al. (2002) provides an example of spreading plots in environmental space. The newly developed Land Environments New Zealand (LENZ, Leathwick et al., in press) is ideal for these purposes. Recording travel time to plots during monitoring would make it possible (after a few years) to develop surfaces of the estimated travel time (or expense) to a location. Within 5–10 years, optimal sampling algorithms could be developed that would use information such as estimated travel time and intensity of interest to provide a surface of optimal sampling intensity. This algorithm could provide the sampling probability surface from which a random probability sampling would choose the sample locations.

5.1.3 Resist the seduction of found data

While past data are useful and should be used where needed, it is important to avoid being seduced into thinking new data are not required. Past data are just that—past. This project aimed to characterise current biodiversity, but much of the data used was badly outdated (Fig. 5). In addition, past data are almost always found data, and have many associated problems, as discussed in Section 1.2.

With continued investment in monitoring, a significant foundation of new data can be developed that is better suited to the intended uses. For a new sampling scheme, it is wiser not to worry about the configuration of existing data, since attempting to correct bias in existing data (much of which in this study are over 20 years old) will result in the new surveys also having a bias. A better approach is to design a new sampling scheme, and use past data as a supplement when it is needed. Within a number of years of monitoring, there will be sufficient information to greatly reduce the need to employ past data.

5.1.4 Permanent plots: positives and negatives

For long-term monitoring, especially of trends, permanent plots have a powerful advantage, and a distinct drawback. Repeated measurements on plots provide a very powerful approach to detecting trends in biodiversity characteristics. The ability of repeat measurement on permanent plots to detect trends has been taken advantage of in the carbon storage monitoring and estimation of Coomes et al. (2001) and Hall et al. (2001). However, permanent, marked plots have a major disadvantage. If characterisation of biodiversity for a region is based on a number of known, permanent plots, the most cost-effective means to improve the 'score' for an area is to manage the plots alone, at the expense of the surrounding areas. An example where the advantages and disadvantages of permanent plots have been recognised is the national monitoring protocol for possum residual trap catch (NPCA 2001). Repeated measurements on the same trap lines are generally avoided to prevent possum trapping and poisoning contractors from focusing on the locations of the trap lines. One potential solution to the problem is to change some proportion (say 25%) of the permanent plots every year (Rob Allen, pers. comm.). This appealing solution would allow remeasurements on many plots, as well as measurements from a significant number of new plots.

5.1.5 Sample both DOC and non-DOC land

Conservation is about private land as well as the land managed by DOC for conservation. Making effective conservation decisions for a particular area also requires information about what is happening in all the surrounding land. Incentives and other programmes for promoting conservation on private land ultimately requires information about the status and trends of biodiversity on this land at regional and national scales. Thus sampling must include plots on non-DOC land.

5.1.6 A suggested sampling scheme

A long-term monitoring scheme should begin from first principles, informed but not constrained by past efforts, with methods and sampling schemes designed to address specified issues (see also Allen 2000 for discussion in a New

Zealand context). Ongoing investment in monitoring should be used to develop this scheme. The sampling design should have randomly chosen or grid sampling locations, and sampling probabilities for each sample should be kept and used in analyses.

An excellent sampling scheme for biodiversity monitoring could be constructed by using following approach, which is similar to that proposed by Bellingham, et al. (2000). A base sample could be evenly spread across the region of interest (e.g. a conservancy, or all of New Zealand). This base sample could be collected using either a systematic grid survey, or a simple random sample. Then an additional random sampling scheme could be carried out to sample more intensively in areas of higher interest. These areas of higher interest may be defined from information gained in the base sampling, from earlier sampling, or from independent sources of information (Fig. 10). For instance, for reporting on conservation outcomes from expenditure on conservation, it may be desirable to sample more intensively (and thereby gain more information) in the areas where conservation operations are occurring, and money is being spent.

There are a number of advantages to such an approach. The base sample across the region would give an overview of the region, and ensure all areas were sampled. The additional samples would allow rare environmental or land-cover combinations to be sampled, or greater sampling intensity in areas identified as of particular interest. The base sample and any additional samples could be combined in analyses, since each sample point would have known sampling probability. Auxiliary samples may measure the same parameters as the base sample, or different methods could be designed to capture different information.

As discussed in this report, an excellent way to define the auxiliary samples is to use probability sampling surfaces. The advantage with this approach is that decisions and compromises made in the process will only affect the efficiency of the sample. The ability of the resulting survey to provide rigorous estimates will not be compromised because each sample will retain its sampling probability. The sampling probability surface and the sample locations would then be included as part of the metadata.

In addition, about 5–10% of the samples should be randomly audited for quality control purposes. The simplest audit is a sample remeasurement by an independent agent. The hard reality is that without these audits, contractors realize, sooner or later, that they do not actually have to measure plots.

5.2 STANDARDISE AND IMPROVE VEGETATION PLOT MEASUREMENTS

It may well be that newly designed monitoring systems use new plot designs. However, a number of improvements on the current recce plots became apparent in this study.

One advantage of using existing methods, such as recce plots, for sampling is that a methodology for measuring the plots exists, and the results from these plots can be compared and combined with past information. One disadvantage

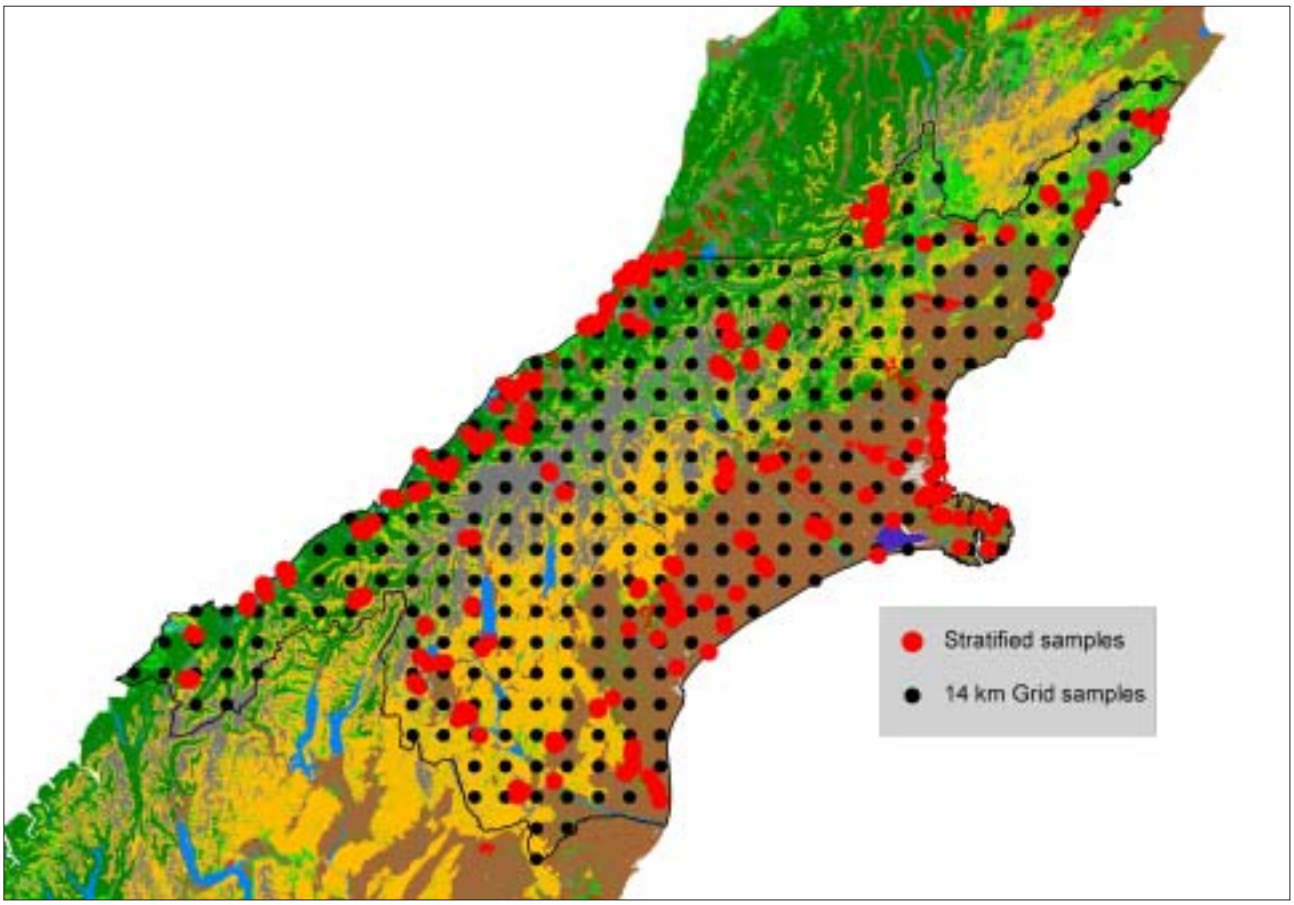

Figure 10. A combined grid and stratified design to sample biodiversity. In this example, a total of abouut 600 sample points is divided equally between a grid sample and a stratified sample. The grid sample is on a 14-km grid and provides an overall sample of the region. The stratified sample is designed to provide more information on rare environments and land covers, and small parcels of conservation land. Samples in the stratified design are clumped to reduce travel time and cost. The relative number of plots in grid and stratified designs can be adjusted to balance the competing constraints of generality and efficiency.

is that there are a number of (to the naïve observer, peculiar) preconceptions as to what the plots represent and how the surveys are carried out. The idea that recce plots characterise a homogeneous vegetation should be abandoned. Recce plots should have defined boundaries and simply record the vegetation at the chosen sample location. Heterogeneity in vegetation is inevitable, and should be accepted. There is no other way to view recce plots when plot locations are chosen prior to carrying out field work. Braun Blanquet cover scores or other schemes should be abandoned and cover should be recorded directly. Cover categories cannot but degrade the accuracy of cover assessments. Under 1%, much resolution is lost with the bottom category of the Braun Blanquet scheme (e.g. in a 20 × 20 m plot, 1% cover is 2 × 2 m and many species are represented by a single individual). In addition, cover categories give a bizarre and artificial distribution of data and require conversion to actual cover for analysis. Most workers familiar with Braun Blanquet resist the change to direct estimation, but eventually find it almost as easy and more precise.

As with all field work, no matter how much attention is paid to detailing the method of implementation, there is a surprising amount of interpretation and subjective decision-making involved in applying methods in the field. Cover

estimates, in particular, have a considerable subjective element. Even the definition of cover has a surprising range of interpretations and needs to be carefully detailed.

As part of the Landcare Research 2000 and 2001 sampling associated with the data used in this project, we have trialled the use of point-height-frequency methods for calibrating recce plots, but these data have not been properly analysed. While these methods have promise, it is clear from the outset that they are not a simple solution. For instance, it is known that the relationship between point-height-frequency methods and cover depends on the grain of the vegetation, and will therefore differ between grasslands and forests.

Other approaches to standardising the recce plots involve assembling the survey crews at the beginning of the season and carrying out repeat measurements with the different crews on the same plots under different vegetation configurations. (Larry Burrows, pers. comm.). This approach has great promise for reducing the variation between crews within a year, but does not help for the next year, unless crews or sampling coordinators are carried over.

5.3 DEVELOP COMBINED BIODIVERSITY INFORMATION SYSTEMS THAT UTILISE A FULL RANGE OF INFORMATION, WITHOUT MIXING AND CONFUSING THE DIFFERENT TYPES

While probability sampling methods have many advantages, they also have shortcomings. Probability surveys are excellent at providing an unbiased and overall characterisation of a region, but they may miss rare events and, generally, cannot incorporate existing knowledge. These shortcomings led Overton et al. (2000, 2002) to develop hybrid methods of biodiversity assessment that combine probability sampling with descriptive surveys. This approach could be generalised to a combined approach for biodiversity assessment that could utilise information of a range of quality and rigor without compromising the better information.

There are several arguments for this sort of system. The sheer range of interests managed by conservation means that no one approach to providing information will suit all purposes. Furthermore, the diversity and number of interests, and the potential expense of monitoring and measuring them all, suggest that, at least for some, the information available will be unreliable or based on expert judgment. There may, however, be value in capturing as much of that information as possible and using it effectively.

The predictions of weed abundance and vegetation produced in this report are based on fairly restrictive and expensive data. The methods use quantitative plot data, and work best to predict the regional patterns of the entities that are best represented in the data, such as the condition indices and the more common weeds and native plants. Restriction to the most reliable and rigorous data is an effective way to get the most rigorous estimates. However, this also results in the avoidance of making predictions for features where there is not solid information. For instance, in this project, we have avoided making

predictions for some weed species for which there was a large component of unidentified congeners in the data. This prevents some of the potential problems of poor predictions for those species, but it does not address the need for the best spatial estimates of those species. If these species are those for which control programs are being carried out, there may be a very real need to make the best prediction possible for those species, keeping in mind that the information is less robust than for species that are reliably identified.

This argument can be taken a step further to consider attributes that are important, but rare across the landscape. Continuing with the weed example, there might be a new weed, which is poorly identified or has only been recorded from scattered locations, which would not have reliable data for the GRASP process demonstrated here. There might, however, be a very real need for spatial estimates of the weed's distribution. At the very minimum, it would be useful to capture data on the locations in which it has been recorded. Note that this process is fraught with difficulties; e.g. many rare weeds may not be identified, others may be over-reported; and some common weeds may be under-reported. Despite all these problems, systematic conservation management will still require information on these sorts of weeds if decisions about their control are to be made objectively and the effects of control are to be quantitatively reported.

There is always the allure of using other sorts of data. For instance, if DOC workers notice a new weed in a new area, it would be very useful to use that information, even if the observation is not part of a formally designed survey. For instance, if Jill Bloggs says she knows there are heavy infestations of this weed at this and that site, it would be useful to have this information included in the decision-making process. But use of this sort of information should not dilute or pollute the predictions made by more quantitative methods and more rigorous data. It would be ideal, therefore to be able to combine different sorts of information to develop an overall picture of the biodiversity condition.

This problem provides exactly the sort of impetus for the hybrid approach to biodiversity assessment used by Overton et al. (2000, 2002) in characterising the biodiversity of roadsides. This approach can be generalised into a staged system that uses information at different degrees of rigour for different purposes. If information for a particular purpose was not available or feasible at a more rigorous level, then a more general level might be used. Furthermore, auxiliary sources of information might be used for different purposes.

While the exact configuration of the system and its components would require development, some examples of levels that might be used are (in decreasing order of rigour):

1. Spatial predictions based upon plot data from rigorously designed probability surveys
2. Spatial predictions based upon existing data
3. Data captured during conservation work
4. Environmental envelopes based on presence-only data and weed observations
5. Descriptive information of sites and areas
6. Expert knowledge and opinion

Spatially explicit predictions of pest abundance suitable for use in MCA and systematic conservation management can be derived from Levels 1, 2, and 3 and informed by level 6. GRASP has considerable advantages for introducing rigour and reproducibility into Levels 1 to 3.

6. Acknowledgments

This report was prepared by Landcare Research for the Department of Conservation as investigation 3418. Theo Stephens has provided continued vision and leadership in developing quantitative conservation frameworks. Mark Smale and Neil Fitzgerald gathered data on the West Coast. Staff at Canterbury Conservancy and at Boffa Miskell designed and executed the sampling in Canterbury. John Leathwick has provided ongoing leadership on species modeling and ecosystem frameworks and reviewed a draft of this report. The NVS database was a significant component of the data used in this modeling, and will serve as the repository for the data developed here. Susan Wiser and Peter Bellingham explained many facets of the NVS database, and Susan Wiser extracted the data used here. Rob Allen engaged in fervent discussions on sampling strategies. Larry Burrows provided useful information on recce plots and ideas for standardising across field crews. M. Anne Austin provided many editorial improvements, David Hunter prepared and revised figures and Rochelle Russell patiently formatted and revised the manuscript.

7. References

Allen, R.B. 1992: RECCE. An inventory method for describing New Zealand vegetation. Forest Research Institute, Christchurch, 25 p .

Allen, R.B. 2000: Forest Health assessment: indicators and methods. Unpublished Landcare report LC9899/111 to Science and Research Unit, Department of Conservation. 38p.

Bellingham, P.J.; Wiser, S.K.; Coomes, D.; Dunningham, A. 2000: Review of permanent plots for long-term monitoring of New Zealand's indigenous forests. *Science for Conservation 151*. 75p.

Cawsey, E.M.; Austin, M.P.; Baker, B.L. 2002: Regional vegetation mapping in Australia: a case study in the practical use of statistical modeling. *Biodiversity and Conservation 11*: 2239–2274.

Chambers, J.M.; Hastie, T.J. 1993: Statistical models in S. Chapman & Hall, London. 608 p.

Coomes, D.A.; Allen, R.B.; Scott, N.A.; Goulding, C.; Beets, P. 2001: Designing systems to monitor carbon stocks in forests and shrublands. *Forest Ecology and Management 5641*: 1–20.

Fielding, A.H.; Bell, J.F. 1997: A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environmental Conservation 24*: 38–49.

Ferrier, S.; Watson, G.; Pearce, J.; Drielsma, M. 2002a. Extended statistical approaches to modelling spatial pattern in biodiversity: the north-east New South Wales experience. I. Species-level modelling. *Biodiversity and Conservation 11*: 2275–2307.

Ferrier, S.; Drielsma, M.; Manion, G.; Watson, G. 2002b: Extended statistical approaches to modelling spatial pattern in biodiversity: the north-east New South Wales experience. II. Community-level modelling. *Biodiversity and Conservation 11: 2309–2338.*

Franklin, J. 1995: Predictive vegetation mapping: geographical modelling of biospatial patterns in relation to environmental gradients. *Progress in Physical Geography 19*: 474–499.

Gustafson, E.J. 1998: Quantifying landscape spatial pattern: What is the state of the art? *Ecosystems 1*: 143–156.

Hall, G.M.J.; Wiser, S.K.; Allen, R.B.; Beets, P.N.; Goulding, C.J. 2001: Strategies to estimate national forest carbon stocks from inventory data: the 1990 New Zealand baseline. *Global Change Biology* 7: 389–403.

Higgins, S.I.; Richardson, D.M. 1998: Pine invasions in the southern hemisphere: modelling interactions between organism, environment and disturbance. *Plant Ecology 135*: 79–93.

Higgins, S.I.; Richardson, D.M.; Cowling, R.M.; Trinder-Smith, T.H. 1999: Predicting the landscape-scale distribution of alien plants and their threat to plant diversity. *Conservation Biology 13(2)*: 303–313.

Leathwick, J.R. 1998: Are New Zealand's *Nothofagus* species in equilibrium with their environment? *Journal of Vegetation Science 9*: 719–732.

Leathwick, J.R. 2001: New Zealand's potential forest pattern as predicted from current species-environment relationships. *New Zealand Journal of Botany 39*: 447–464.

Leathwick, J.R.; Austin, M.P. 2001: Competitive interactions between tree species in New Zealand's old growth indigenous forests. *Ecology 82*: 2560–2573.

Leathwick, J.R.; Whitehead, D. 2001: Soil and atmospheric water deficits, and the distributions of New Zealand's indigenous tree species. *Functional Ecology 15*: 233–242.

Leathwick, J.R.; Wilson, G.H.; Rutledge, D.; Wardle, P.; Morgan, F.J.; Mcleod, M. Kirkpatrick (in press). Land Environments of New Zealand. Bateman.

Lehmann, A.; Leathwick, J.R.; Overton, J.McC. 2002a. Assessing New Zealand fern diversity from spatial predictions of species assemblages. *Biodiversity and Conservation 11*: 2217–2238.

Lehmann, A.; Overton, J.McC.; Leathwick, J.R. 2002b: GRASP: Generalized Regression Analysis and Spatial Prediction. *Ecological Modeling 157*: 189–207.

Margules, C.R.; Pressey, R.L. 2000: Systematic conservation planning. *Nature 405*: 243–253.

NPCA 2001. Protocol for Designers – for possum population monitoring. National Possum Control Agencies, Wellington. 27 p.

Overton, J.McC.; Young, T.C.; Overton, W.S. 1993: Using 'found' data to augment a probability sample: procedure and case study. *Environmental Monitoring and Assessment 26*: 65–83.

Overton, J.McC.; Smale, M.C.; Whaley, K.J.; Clarkson, B.D.; Emmett, D.K.; Simcock, R.C. 2000: Measuring and enhancing roadside biodiversity. *Transfund New Zealand Research Report No. 157*. Wellington.

Overton, J.McC.; Smale, M.C.; Whaley, K.J.; Fitzgerald, N.B.; McGlone, H.M. 2002: A Methodology for assessing the biodiversity of road networks: a case Study of the Waikato Region, New Zealand. *Transfund New Zealand Research Report No. 221.* Wellington.

Overton, J.McC.; Stephens, R.T.T.; Leathwick, J.R.; Lehmann, A. 2002: Information pyramids for informed biodiversity conservation. *Biodiversity and Conservation 11*: 2093–2116.

Pearce, J.; Ferrier, S. 2001: The practical value of modelling relative abundance of species for regional conservation planning: a case study. *Biological Conservation 98*: 33–43.

Stephens, R.T.T. 1999: Measuring conservation achievement. Pp. 13–39 in Blaschke, P.M.; Green, K. (Eds): Biodiversity Now! Joint societies conference, Wellington, 1997. Selected papers. Department of Conservation, Wellington, New Zealand.

Stephens, R.T.T.; Brown, D.J.; Thornley, N.J. 2002: Measuring Conservation Achievement: concepts and their application over the Twizel area. *Science for Conservation 200*. 114 p.

Welsh, A.H.; Cunningham, R.B.; Donnelly, C.F.; Lindenmayer, D.B. 1996: Modelling the abundance of rare species: statistical models for counts with extra zeros. *Ecological Modelling 88*: 297–308.

Wiser, S.K.; Rose, A.B. 1997: Two permanent plot methods for monitoring changes in grasslands: a field manual. Manaaki Whenua – Landcare Research, Lincoln, New Zealand.

Wiser, S.K.; Allen, R.B.; Clinton, P.W.; Platt, K.H. 1998: Community structure and forest invasion by an exotic herb over 23 years. *Ecology 79(6)*: 2071–2081.

Wiser, S.K.; Bellingham, P.J.; Burrows, L.E. 2001: Managing biodiversity information: development of New Zealand's National vegetation Survey databank. *New Zealand Journal of Ecology 25*: 1–17.

Appendix 1: Sample predictions and analyses

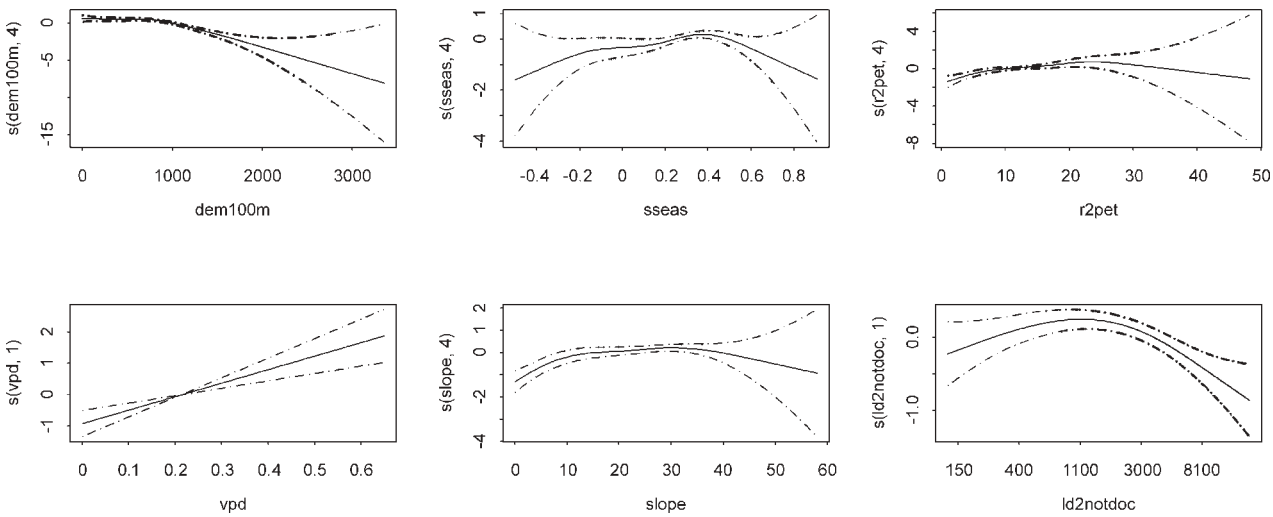

Figure A1. Models of plot sampling density for DOC land.

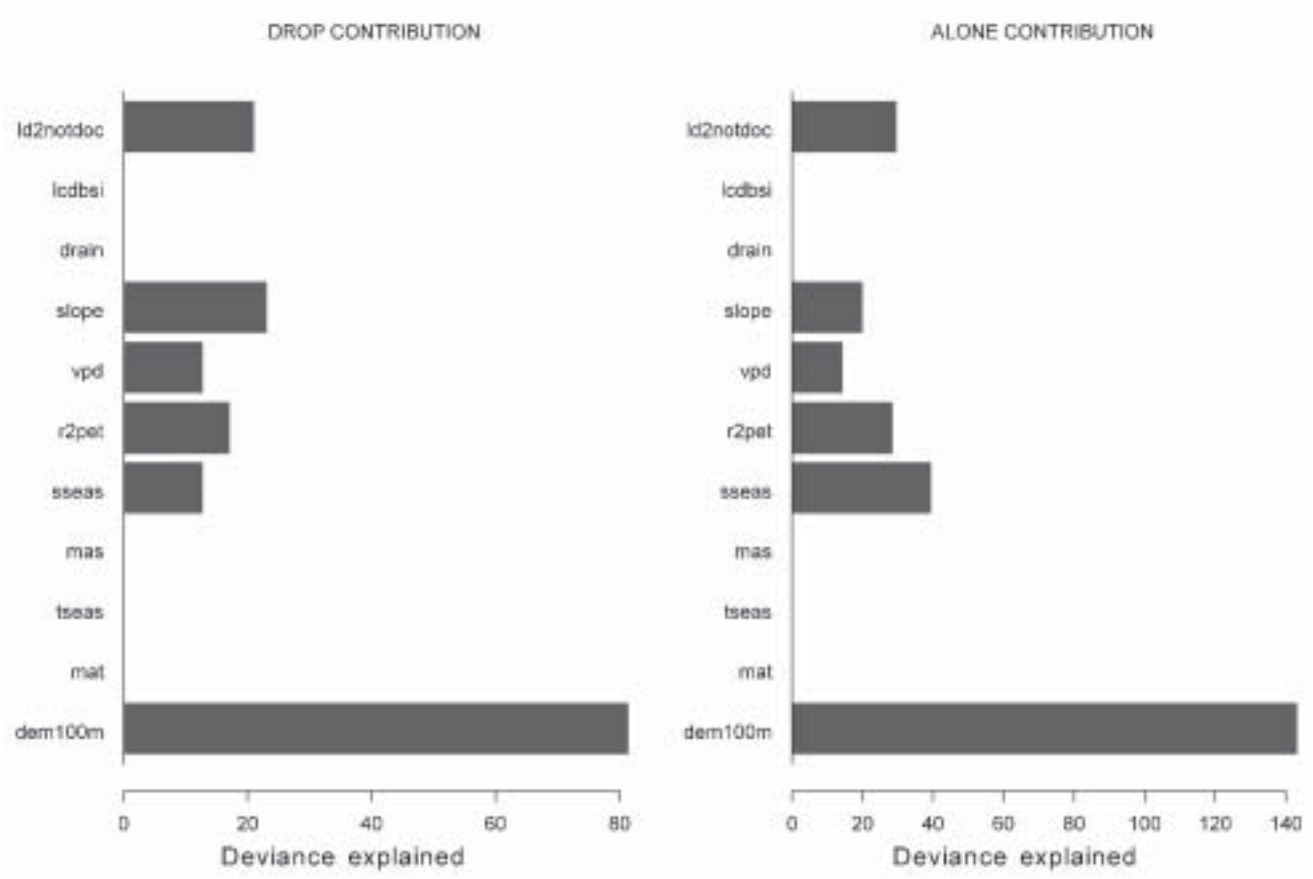

Figure A2. Contributions to the models of plot sampling density for DOC land. The alone contribution for a predictor variable is defined as the amount of deviance explained when that predictor is alone in the model. Drop contributions are defined as the drop in deviance when the predictor is removed from the full, final model.

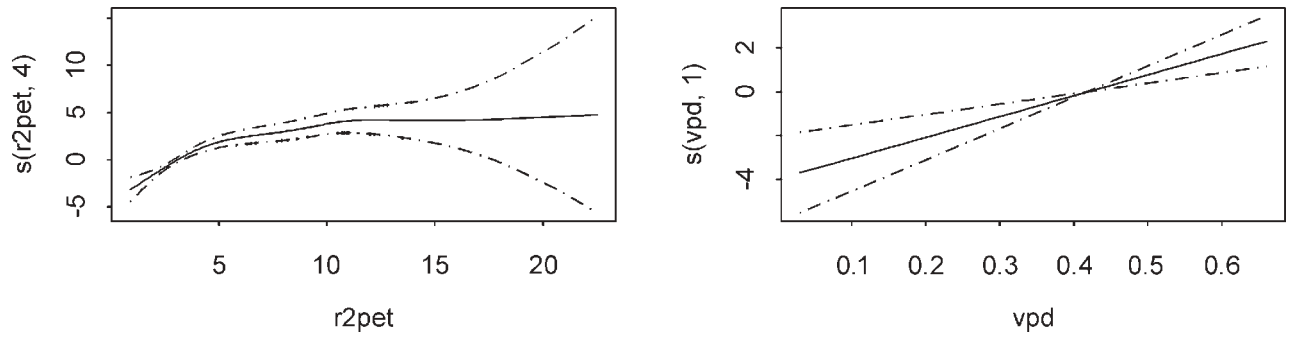

Figure A3. Models of plot sampling density for non-DOC land.

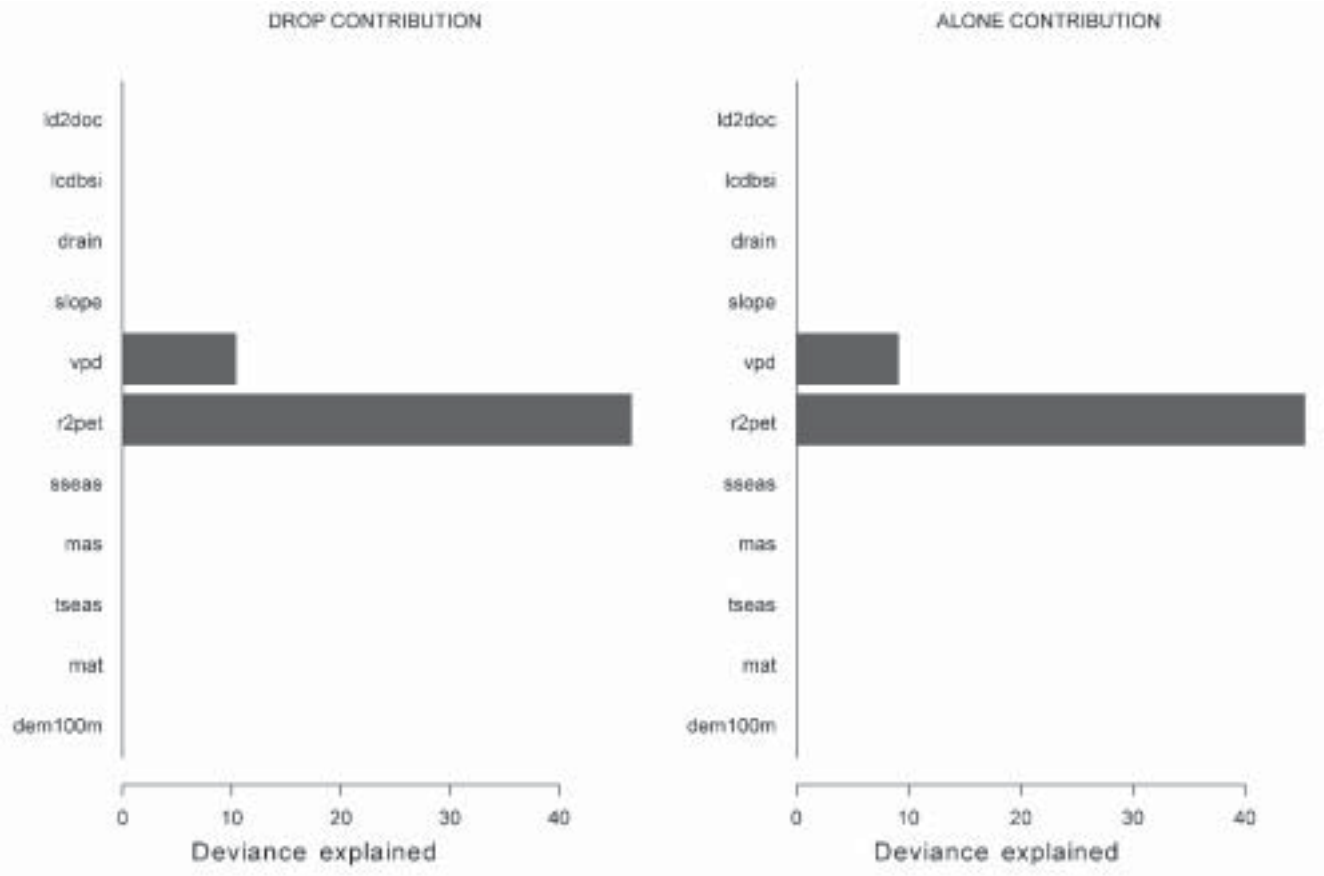

Figure A4. Contributions to the models of plot sampling density for non-DOC land.

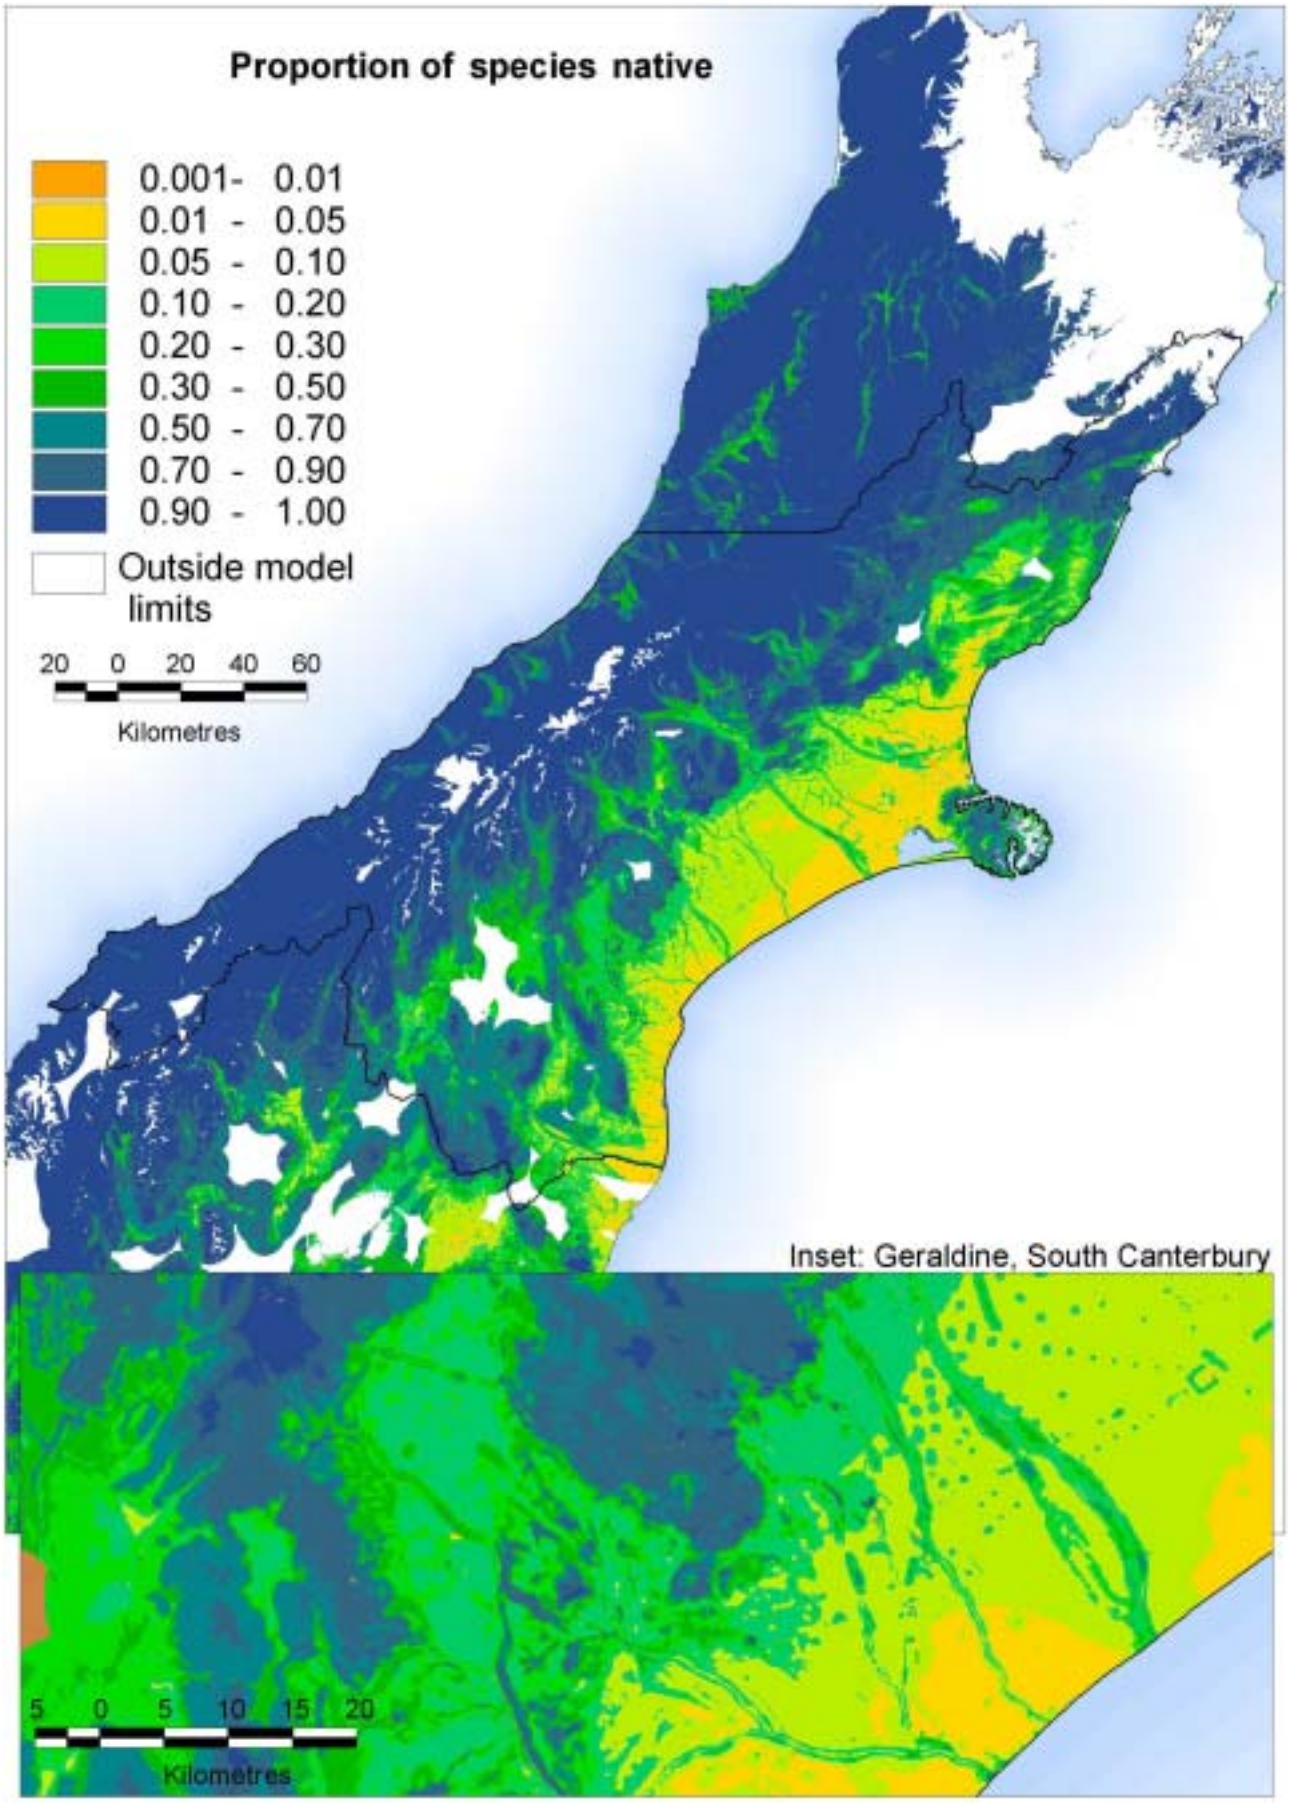

Figure A5. Prediction of the proportion of species native.

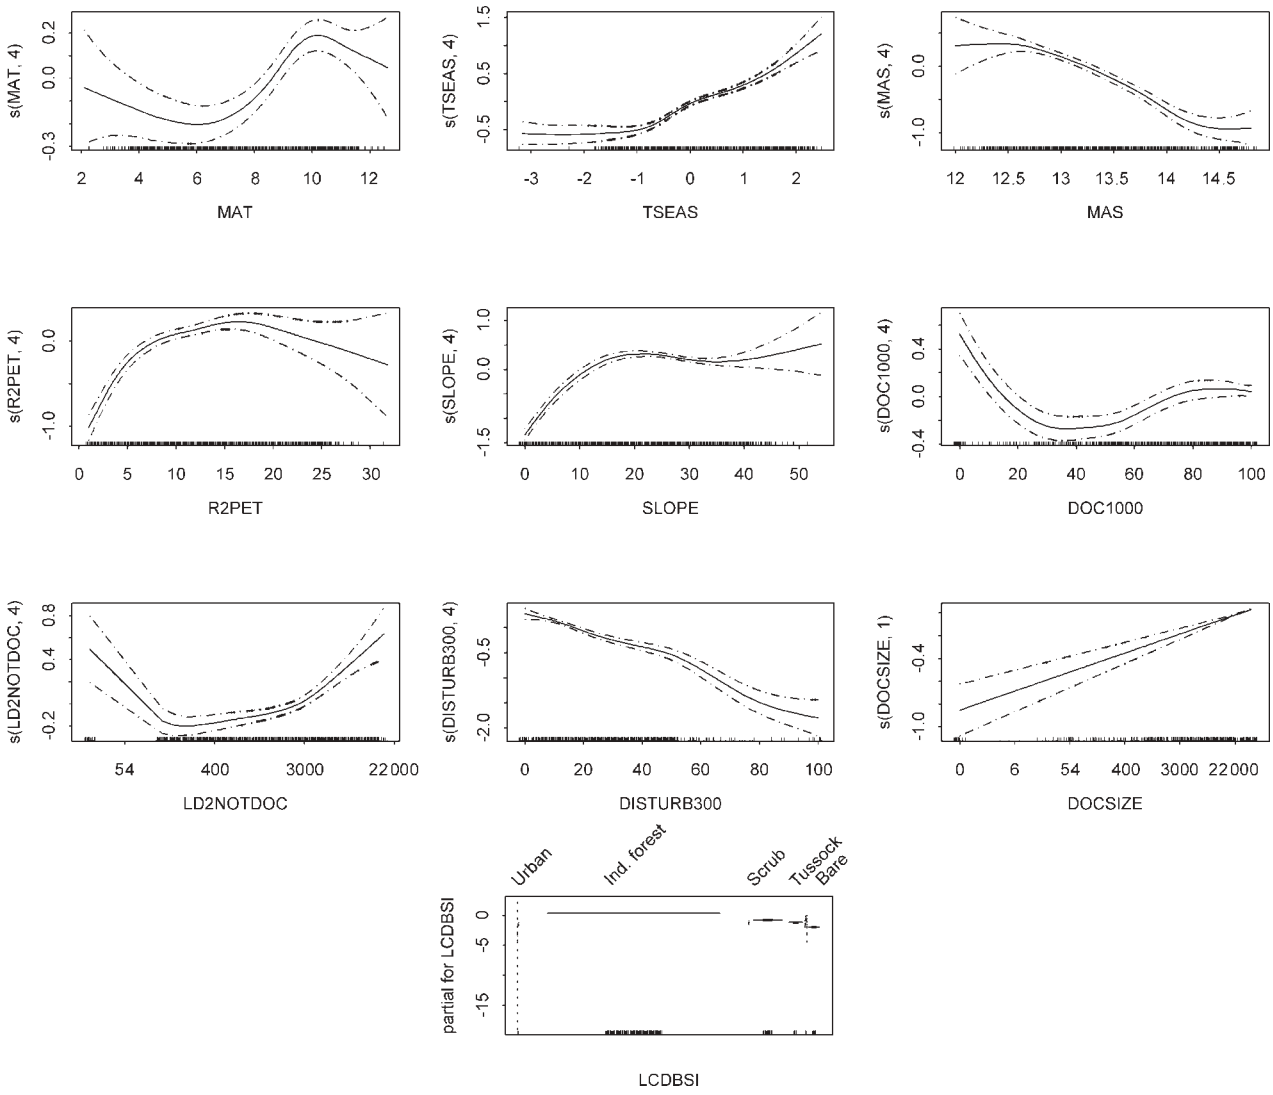

Figure A6. Model of the proportion of species native for DOC land.

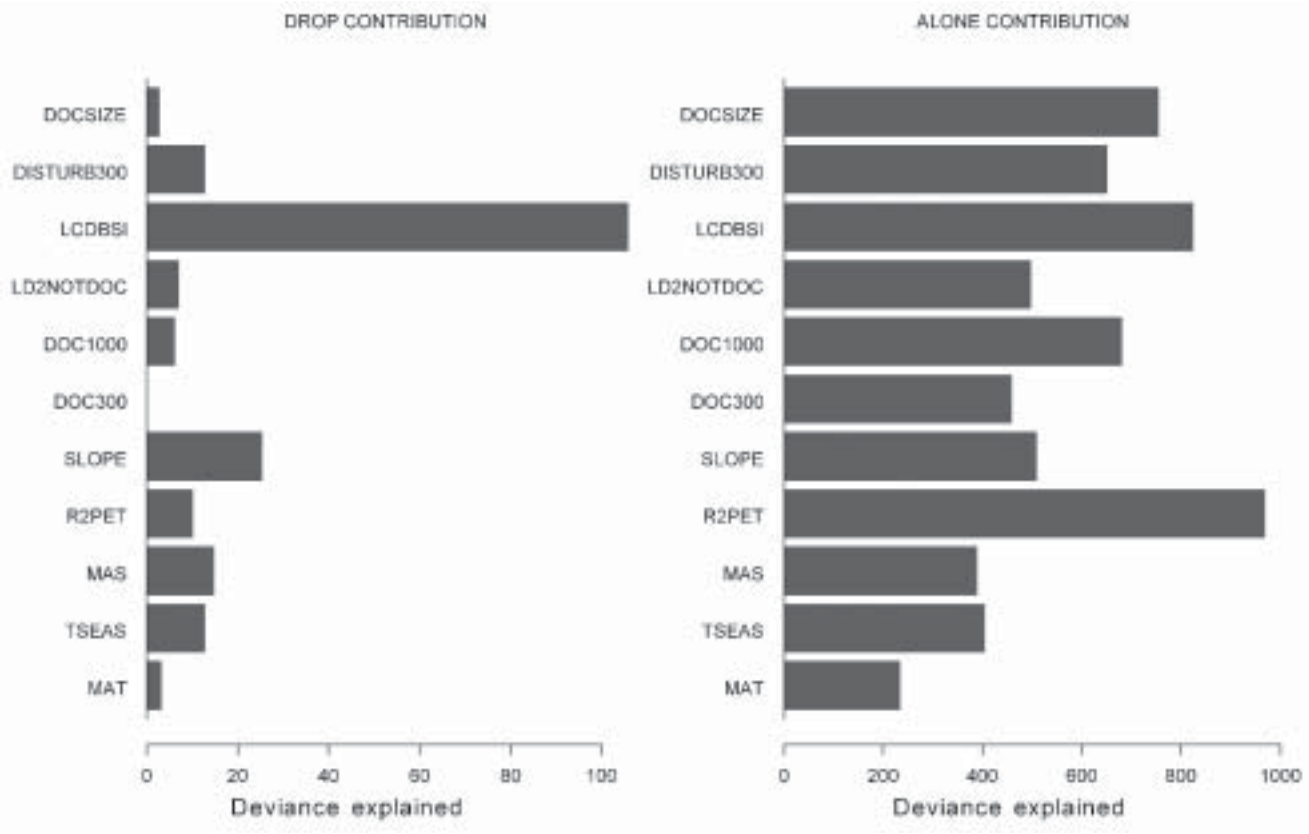

Figure A7. Contributions for the model of the proportion of species native for DOC land.

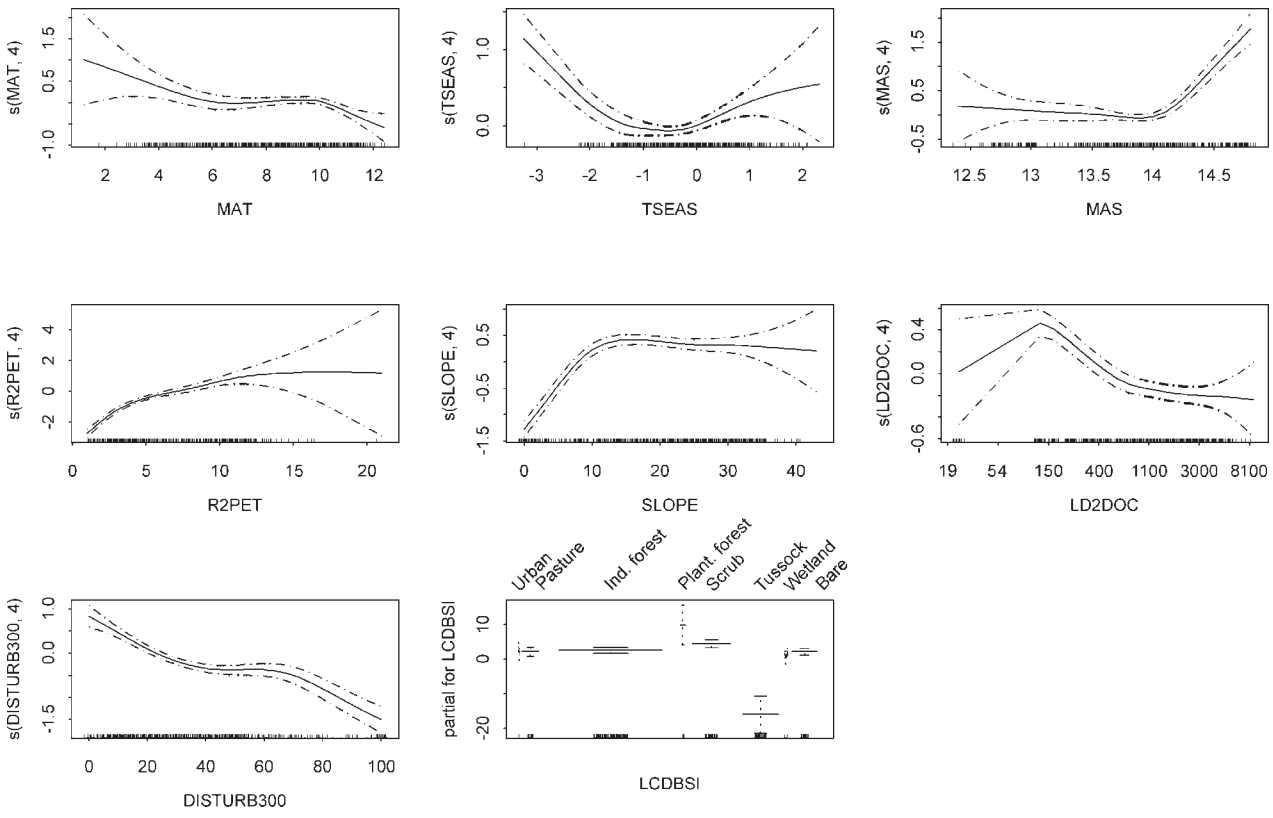

Figure A8. Model of the proportion of species native for non-DOC land.

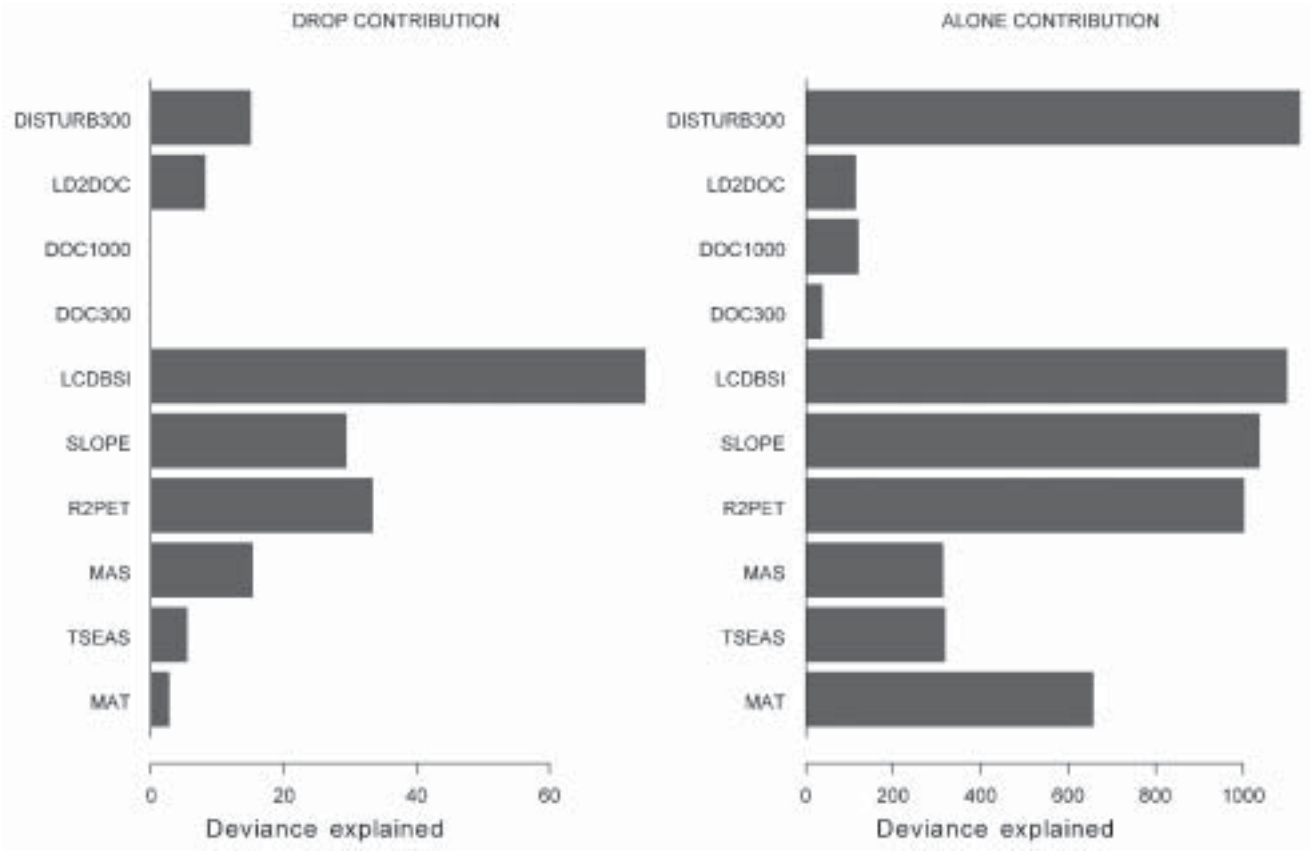

Figure A9. Contributions for the model of the proportion of species native for non-DOC land.

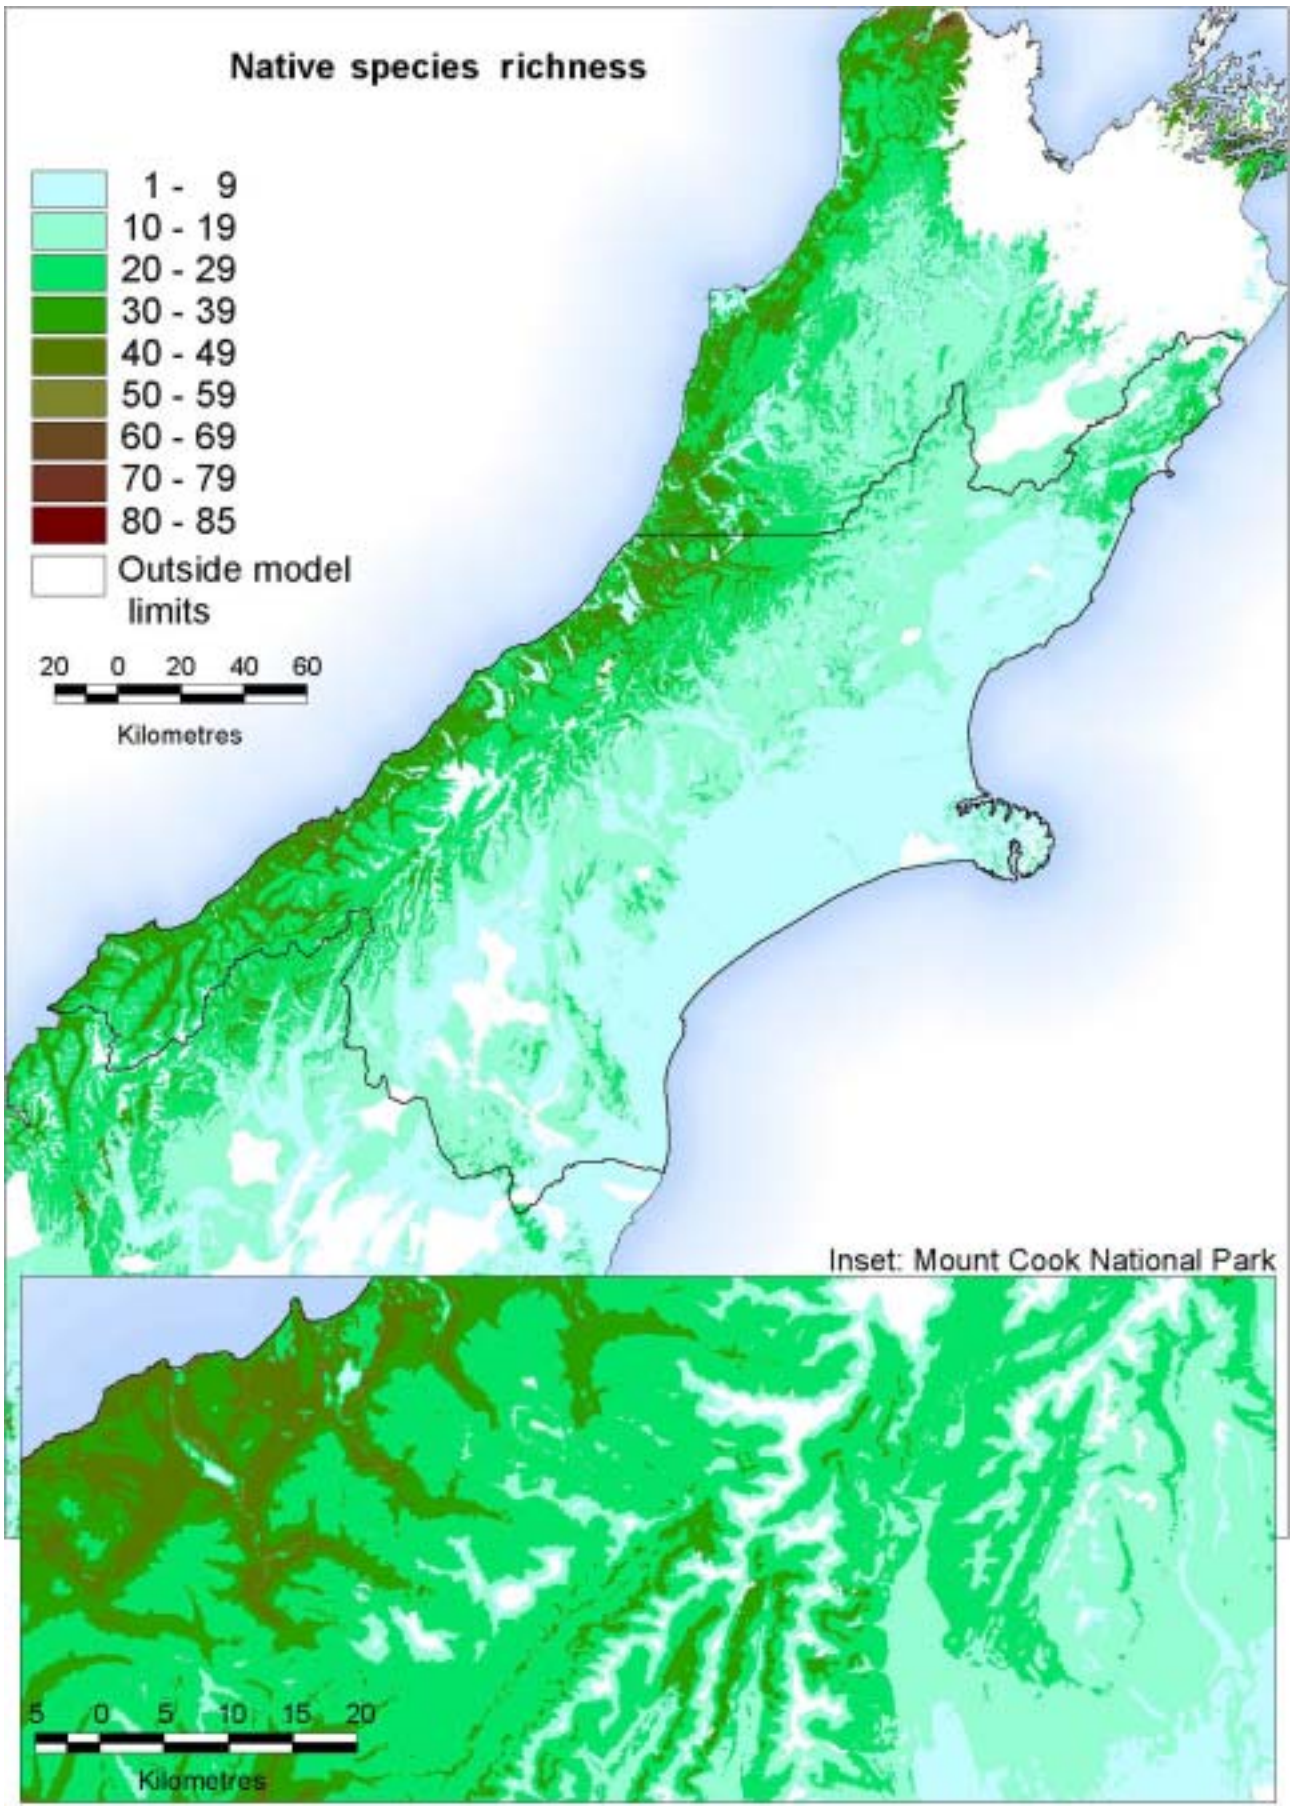

Figure A10. Prediction of native species richness.

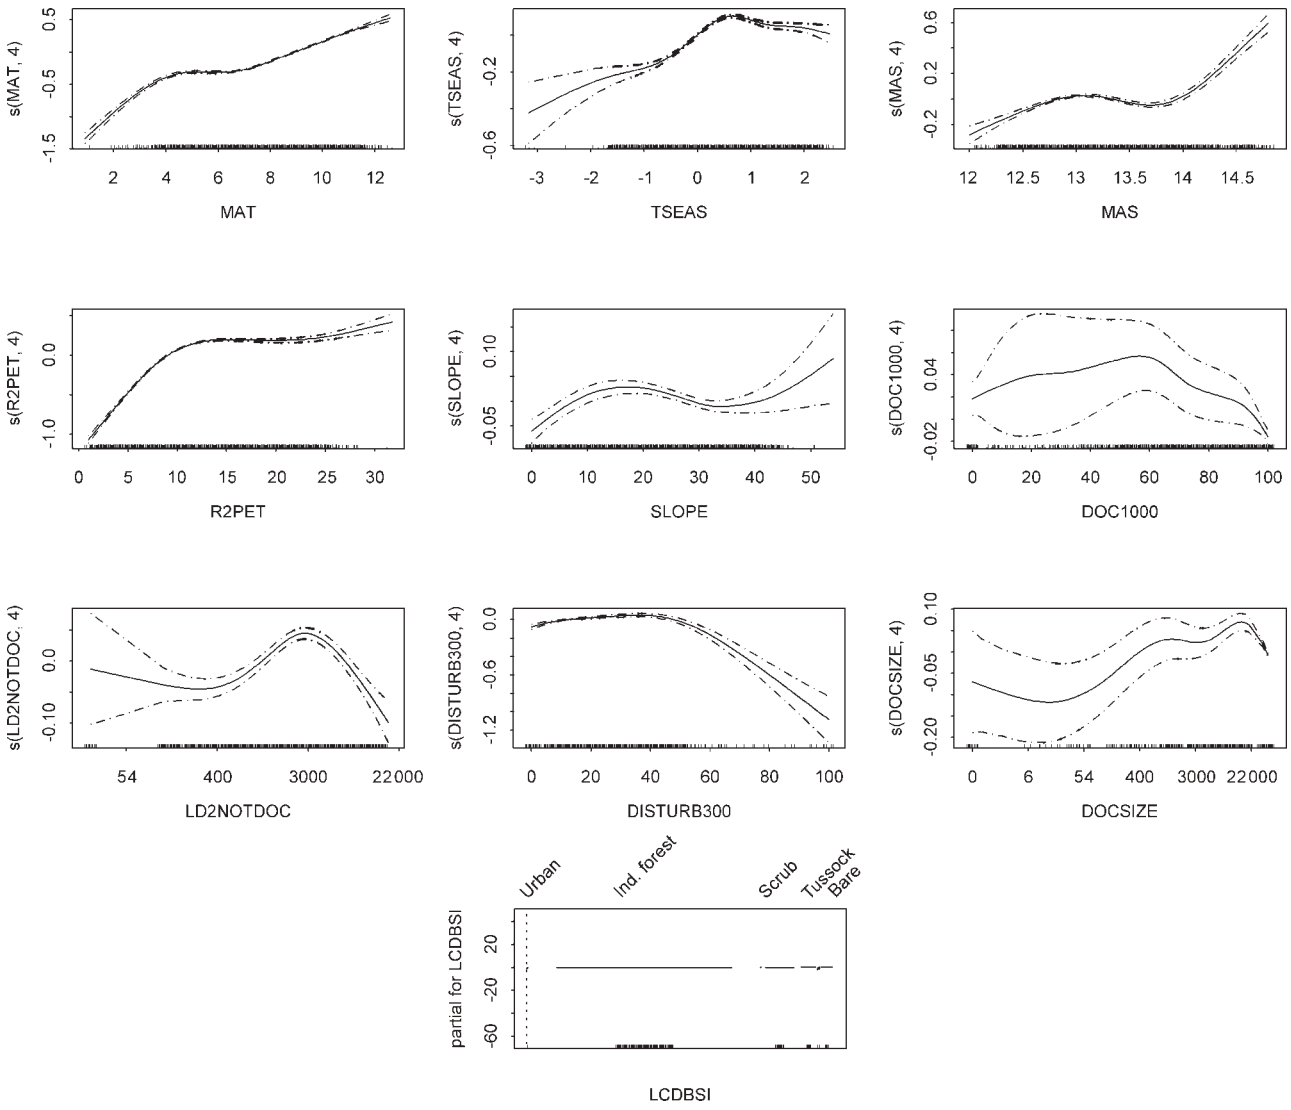

Figure A11. Model of native species richness for DOC land.

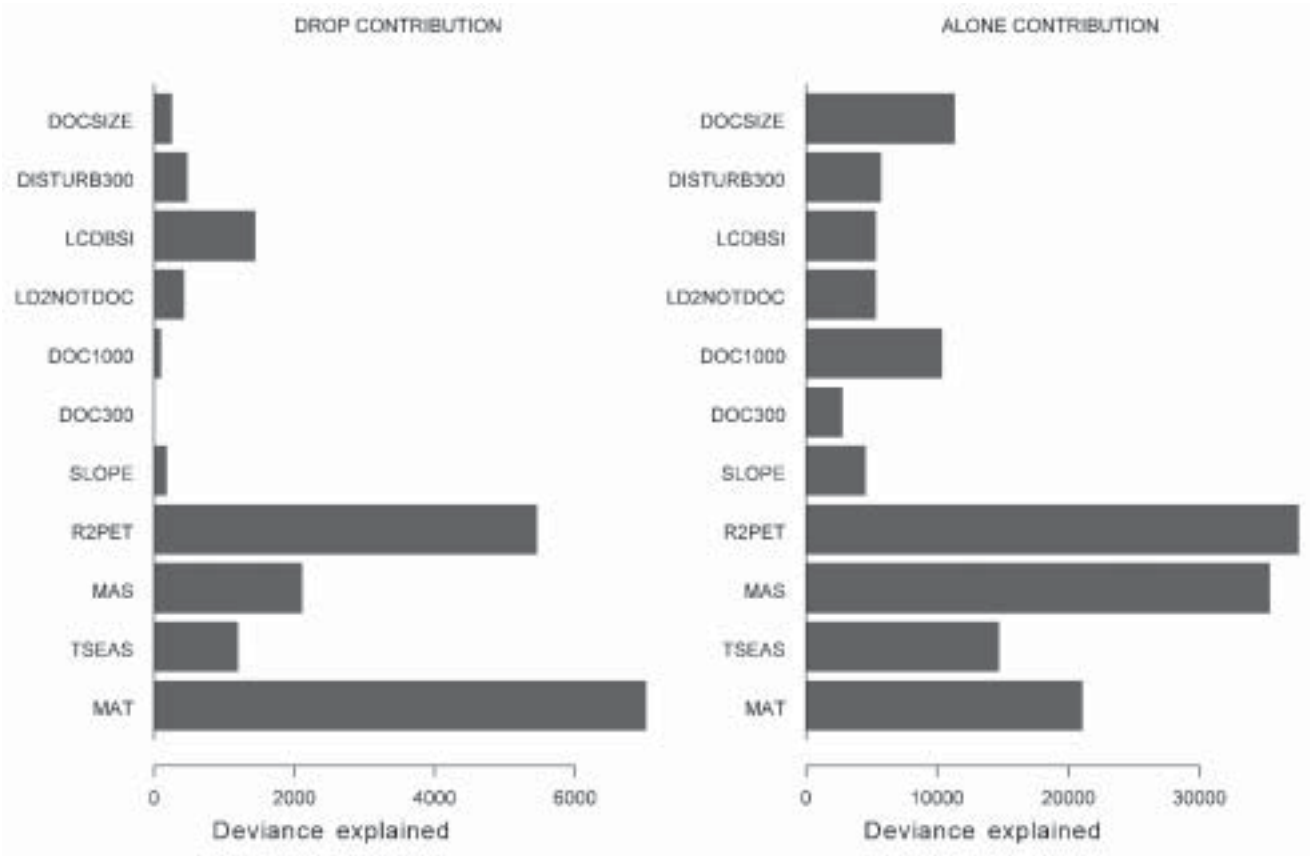

Figure A12. Contributions for the model of native species richness for DOC land.

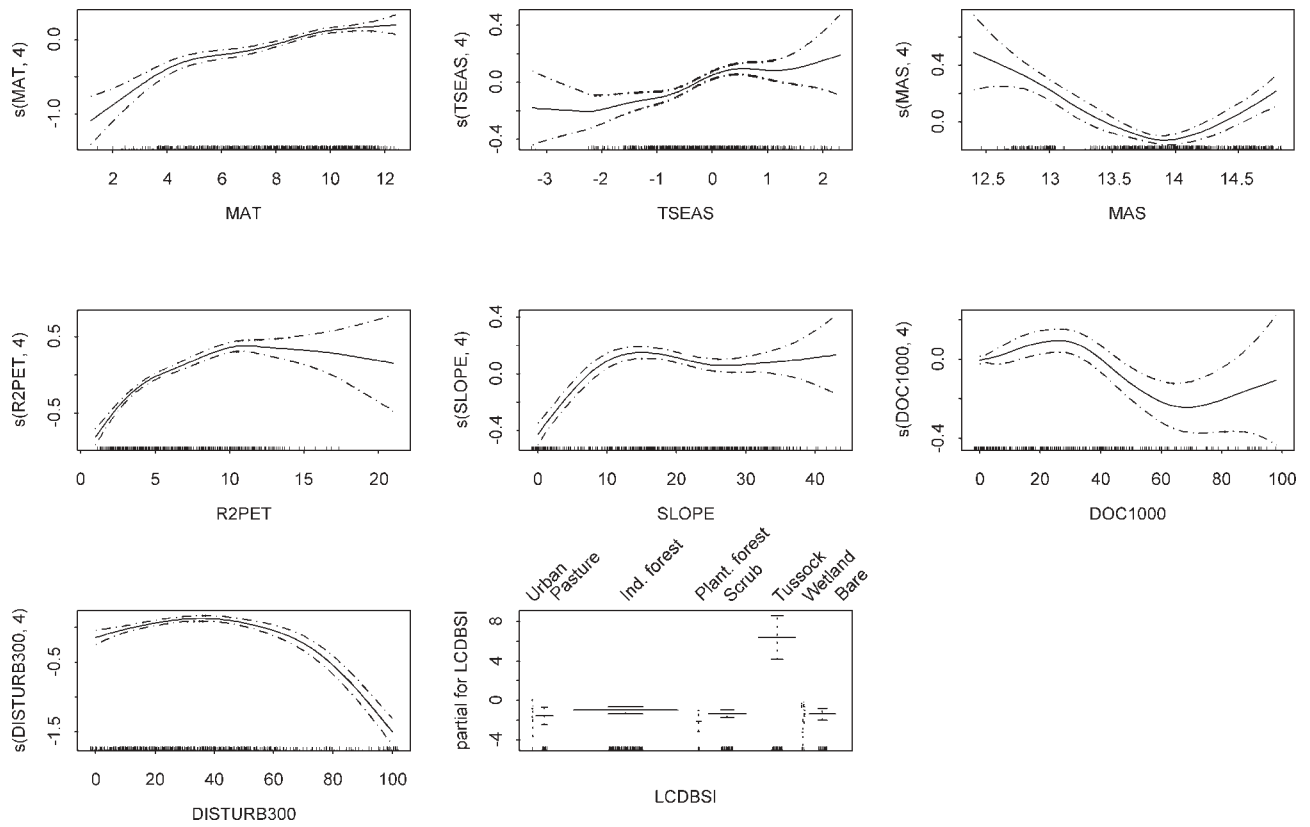

Figure A13. Model of native species richness for non-DOC land.

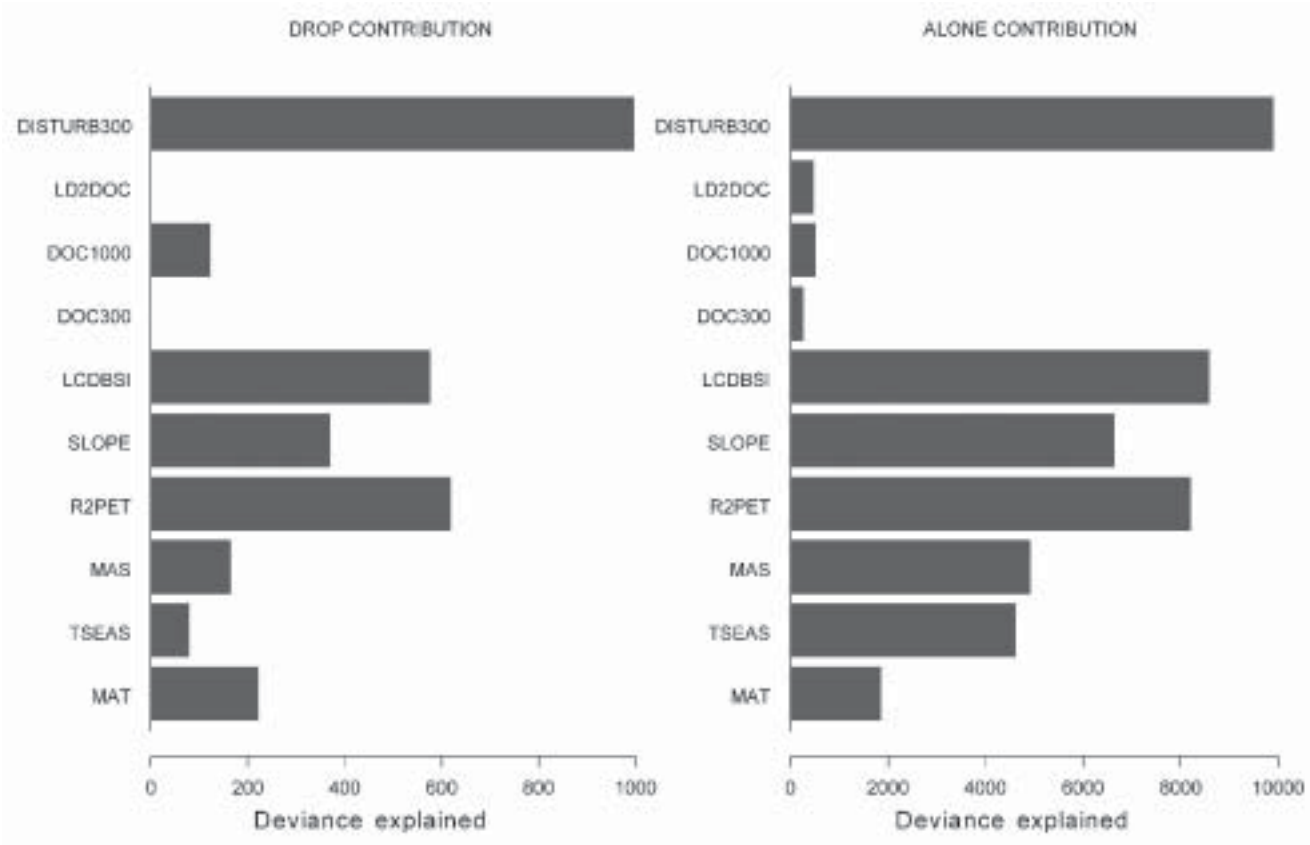

Figure A14. Contributions for the model of native species richness for non-DOC land.

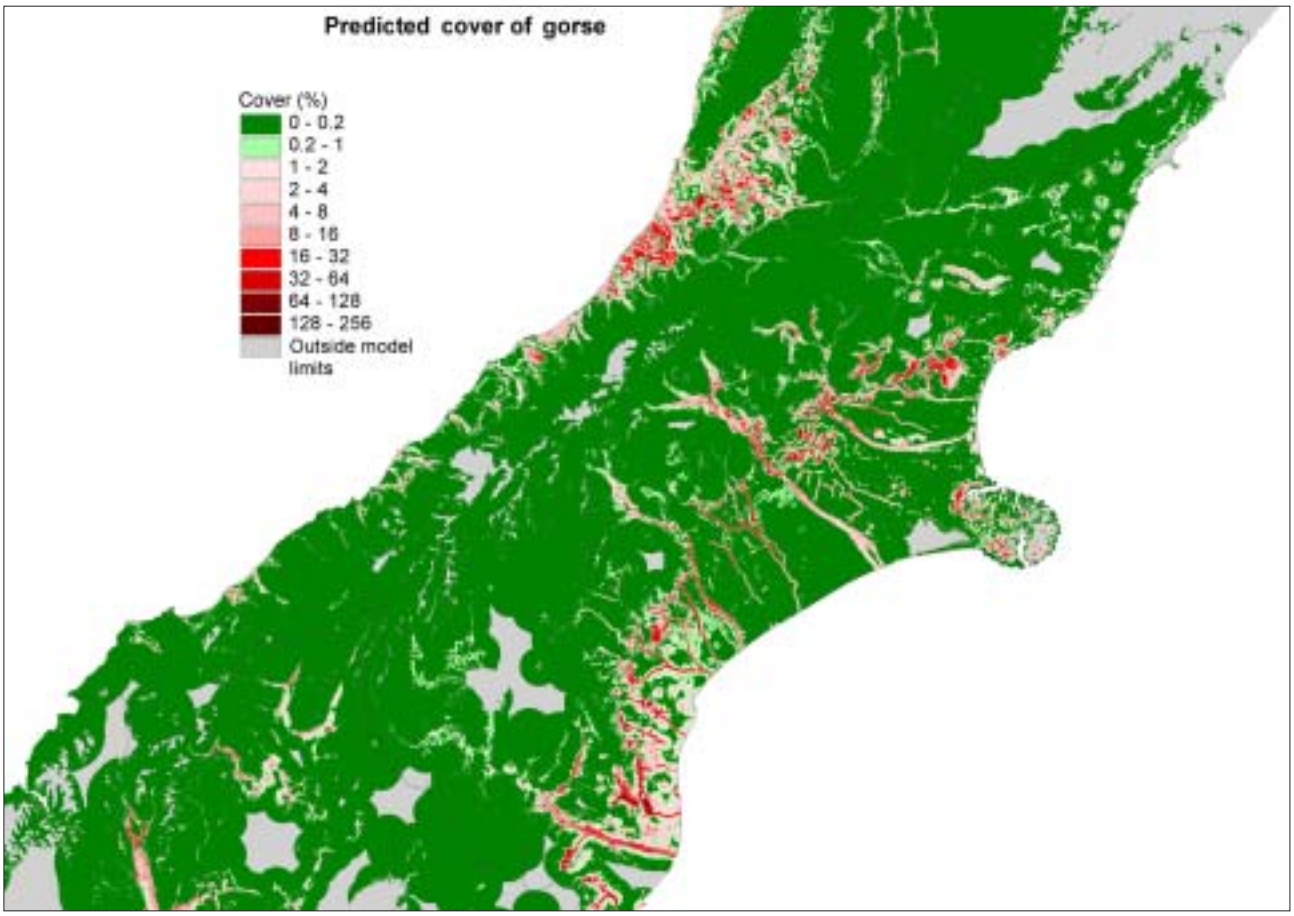

Figure A15. Prediction of the abundance of gorse, *Ulex europea*.

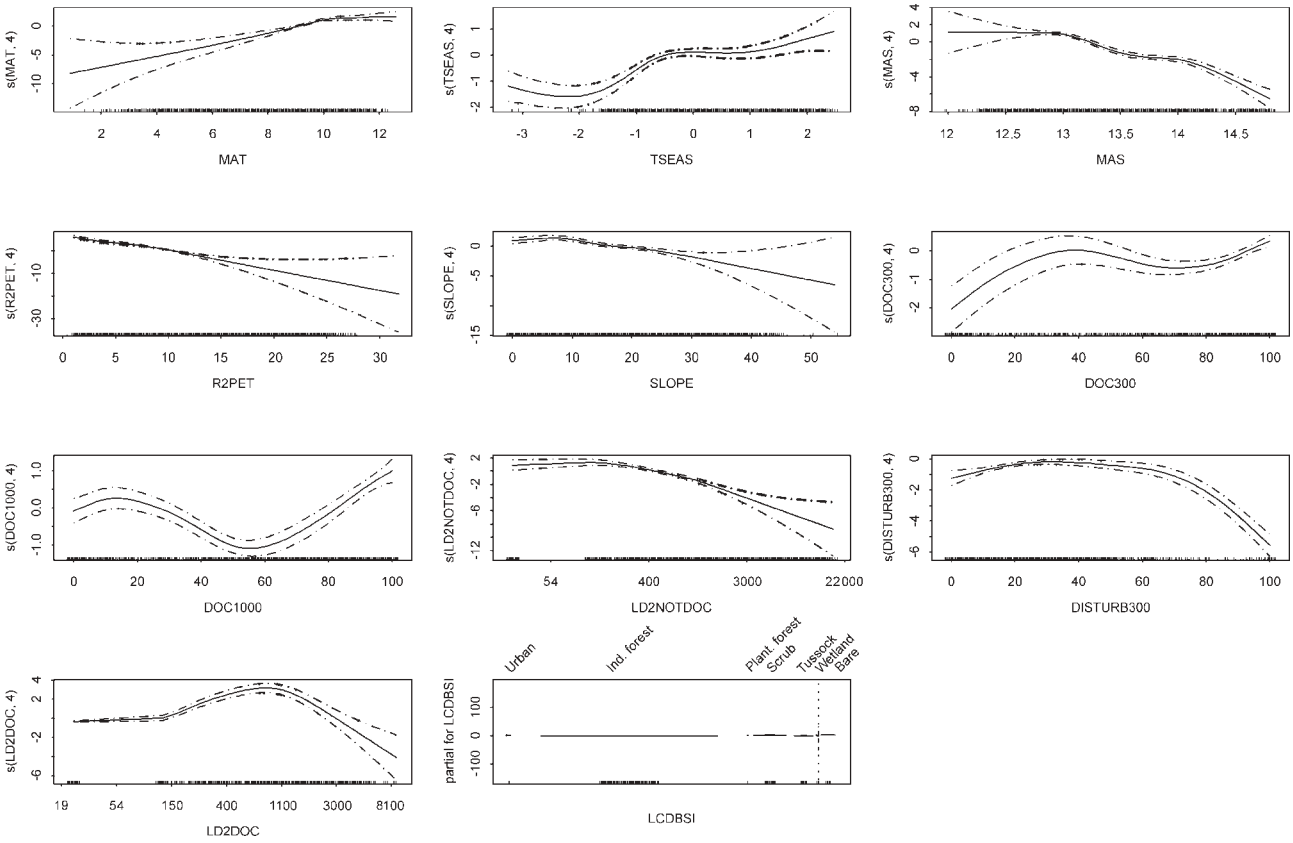

Figure A16. Model of the presence-absence of gorse for DOC and non-DOC land.

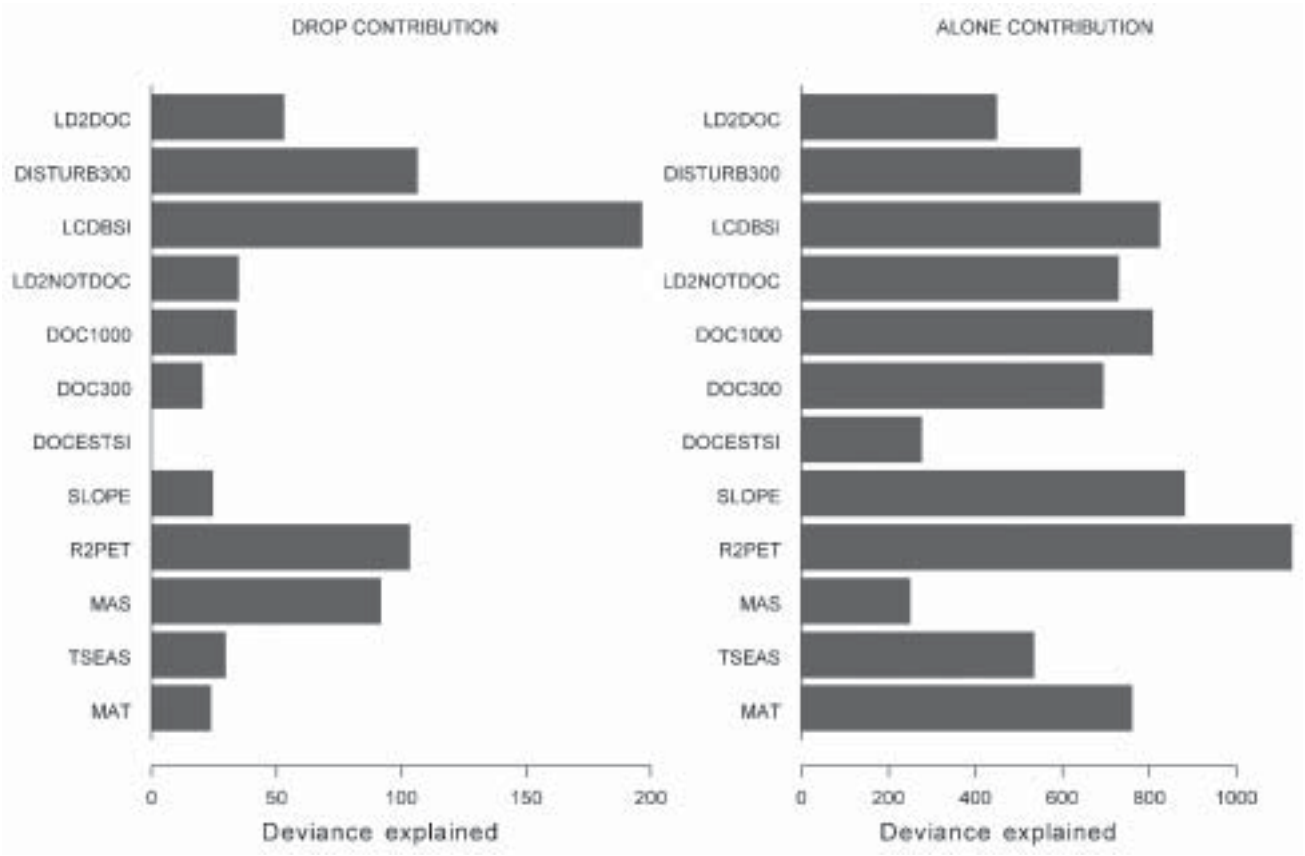

Figure A17. Contributions for the model of presence-absence of gorse.

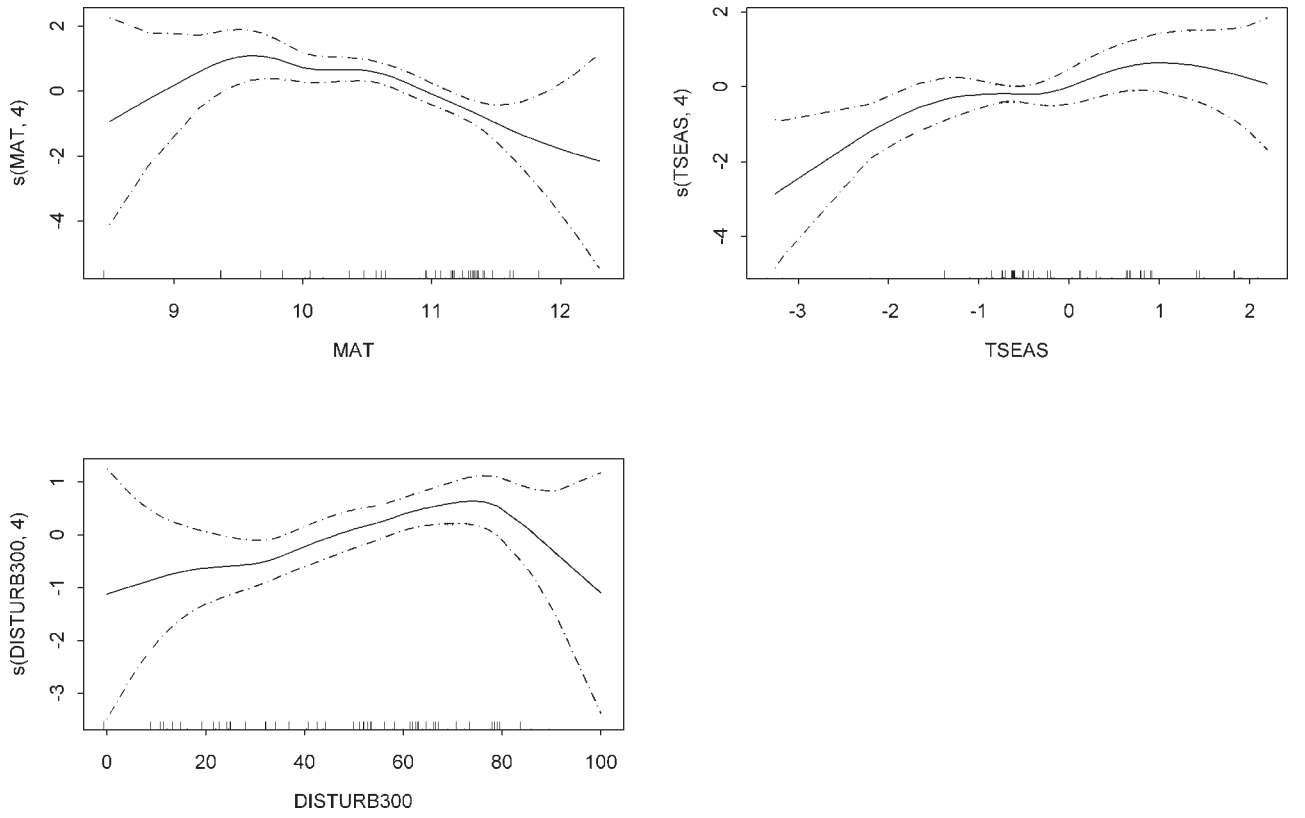

Figure A18. Model of gorse abundance, performed over all sites where gorse was present. Plot methods included recording cover scores.

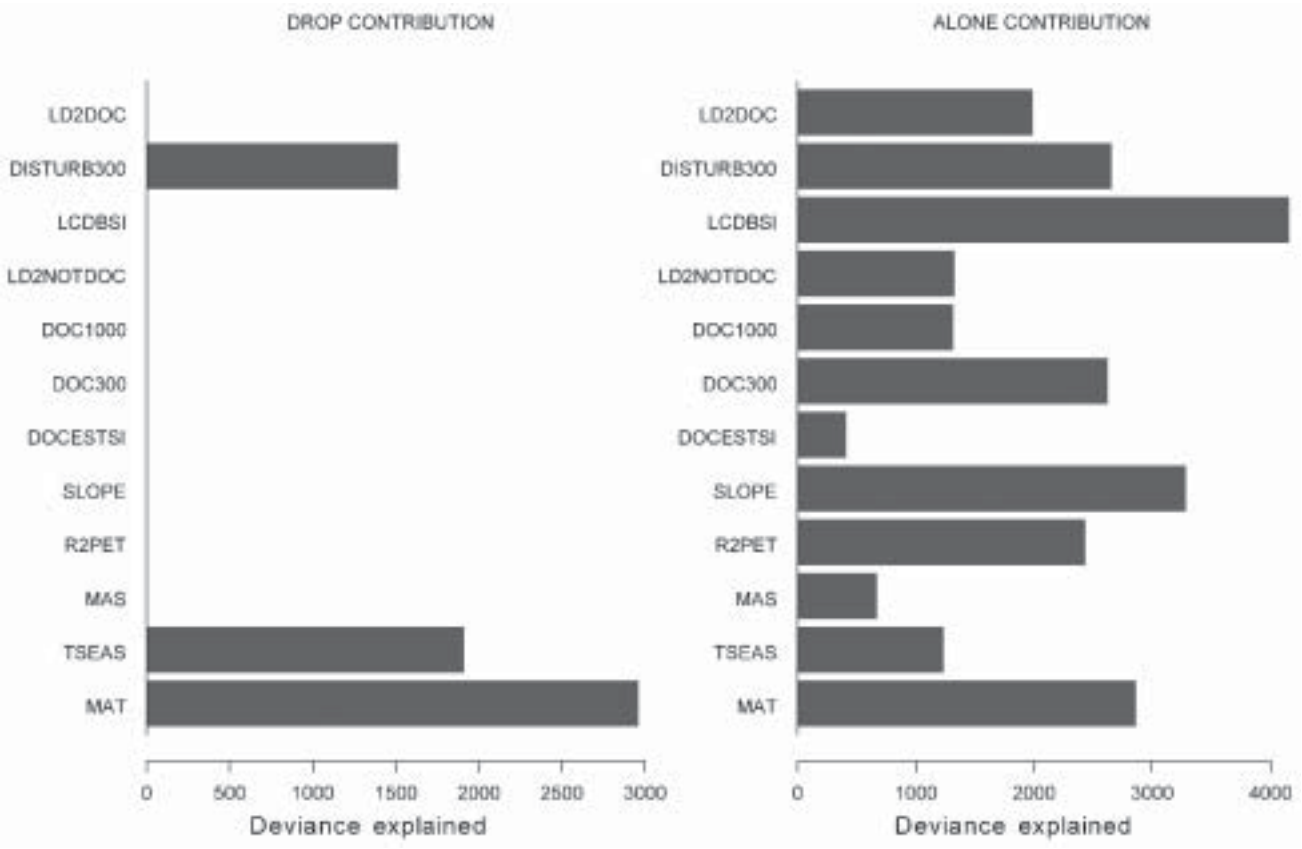

Figure A19. Contributions for the model of gorse abundance for sites where it was present.

Glossary

Alone/drop contributions

These are two different ways of looking at the importance, to the model, of each spatial predictor variable. The alone contribution of a spatial predictor variable shows the deviance explained by a model with only that variable in the model. The alone contribution is calculated by having no variables in the model, putting each variable into the model by itself, and recording the amount of deviance explained. The drop contribution of a spatial predictor variable is the difference in explained deviance between the full final model and a model with the variable excluded. The drop contribution is calculated by taking the final model, and dropping out single variables, one at a time, and recording the decrease in the amount of deviance explained. If the variable in question is not in the final model, the drop contribution is defined as zero. The drop contribution of a predictor variable is often much smaller than the alone contribution, because of the redundancy between predictor variables.

Backwards/forwards stepwise selection

The choice of variables to include in multiple regression models is often made by stepwise selection, using a set of criteria to determine whether predictor variables are included in the model. In forwards stepwise selection, the most important variable is added first, followed by the next most important variable, step by step, until all variables that satisfy the criteria are included in the model. Backwards stepwise selection begins with all the variables in the model, and drops out variables that do not satisfy the criteria, starting with the weakest variable, step by step, until only the variables that satisfy the criteria are included in the model.

Generalised information

Generalised information is derived from more detailed underlying information or data. Examples of generalised information include theories or models which may capture the overall or average behavior of a system, and classifications such as vegetation classes and environmental and biotic domains that generalise the pattern of many underlying species into broad classes of biotic composition. GRASP is an important method for generalising information. See also Information pyramid.

Inclusion probabilities

Inclusion probabilities are a crucial part of probability samples. The inclusion probabilities give the probability that any element of the target population is included in the sample. For a simple random sample, the inclusion probabilities are all equal. For stratified designs, some parts of the target population have higher probabilities than others. The inclusion probabilities are used to weight the observations in analyses, and are thus an important part of metadata.

Information pyramid

Information pyramids represent a paradigm for informed ecosystem management (Overton et al. 2002). Raw data forms the base of the pyramid. This data is integrated with other forms of data and generalised to produce the derived information required for informing ecosystem management. Results at one level are used as input for higher levels of integration and generalisation.

Integrated information

Oftentimes, the process of generalising information requires that different sorts of information be integrated. In this report, we integrate biotic information with environmental information to produce spatial predictions. See also information pyramid.

Masked/masking

For spatial predictions, a mask identifies areas for which the prediction is considered invalid or unreliable. In this project, predictions are masked to avoid predicting outside the range of the data available for the model.

Metadata

Metadata is data that describes data. Often, this can consist of information on the people that collected or entered the data, or the methods used to generate the data. Metadata should also include information on the sampling design used to collect the data.

Probability sampling

Sampling designs in which the locations or elements to sample are chosen according to a random sampling scheme. Examples of probability sampling include simple random sampling, stratified random sampling and most grid sampling. The use of probability sampling can ensure that the results are rigorous and unbiased.

Regression modelling/analysis

An analysis that produces statistical relationships between a response variable and one or more predictor variables. Examples of regression modeling include simple linear regression, generalised linear models (GLMs), and generalised additive models (GAMs). A multiple regression model regresses a single response variable against a number of predictor variables.

Spatial prediction

A paper or digital map that gives predicted values across a region. In this study, the spatial predictions were made in a GIS using the GRASP models, and give the predicted values for every 100 m pixel of the study region, except for the pixels that are in the areas that are masked out.

Spatial predictors

See variables, spatial and variables, predictor.

Spectral characteristics

Used in this report as a generic way of referring to remote sensing information. This is often reflectances in different wavelength bands, with values for each of the pixels in the image.

Splus

A statistical package, with a programming language, for advanced statistical computing (Chambers & Hastie 1993).

Surface

In this study, surfaces give a predicted value of a variable across geographic space, and are stored in a GIS. Surfaces consist of many pixels, with a value of the surface for each pixel. For instance, the surface of mean annual temperature used here gives a predicted value of temperature for every 100 m (1 ha) pixel of the study region. Surfaces are also spatial predictions, because the variable has not been measured for every pixel of the surface. However, they are spatial predictions that were produced at lower levels of the information pyramid, and hence become spatial predictor variables for the current use.

Variables

Cadastral—Spatial variables that show land ownership.

Categorical—Variables that have discrete categories, or classes, such as land cover.

Continuous—Variables that vary continuously, such as temperature.

Fragmentation—Variables that capture the patchy distribution of habitats across the landscape, and the influences of humans on landscape pattern. Patch-based indices of fragmentation focus on 'patches' of habitat, surrounded by unsuitable habitat. Continuous measures of fragmentation developed in this report deal better with the continuous nature of habitat and give a characterisation of fragmentation that varies continuously across the landscape.

Response—In a regression model, or spatial prediction, this is the variable that is being predicted. In the case of this report, response variables include the condition indices and the weed presence/absences or abundances

Predictor—Variables that are used to predict the response variables. They are the variables on the x-axis of the regression models. Oftentimes, predictor variables were response variables at lower levels of the information pyramid, and become predictors at higher levels.

Spatial—Variables for which a surface exists in a GIS. Spatial predictor variables include the variables listed in Table 2. Response variables also become spatial variables once spatial predictions are made of them.