Global Ecology and Biogeography, (Global Ecol. Biogeogr.) (2017) **26,** 243–258

Global Ecology and Biogeography A Journal of Macroecology

MACROECOLOGICAL METHODS

Integrating occurrence data and expert maps for improved species range predictions

Cory Merow[1,2]*, Adam M. Wilson[1,3] and Walter Jetz[1,4]

[1]Ecology and Evolutionary Biology, Yale University, 165 Prospect Street, New Haven, CT, USA, [2]Ecology and Evolutionary Biology, University of Connecticut, Storrs, CT, USA, [3]Department of Geography, University at Buffalo, Wilkeson Quad, Buffalo, NY, USA, [4]Department of Life Sciences, Imperial College London Silwood Park Campus Ascot, Berkshire SL5 7PY, UK

ABSTRACT

Aim Knowledge of species geographical distributions is critical for many ecological and evolutionary questions and underpins effective conservation decision-making, yet it is usually limited in spatial resolution or reliability. Over large spatial extents, range predictions are typically derived from expert knowledge or, increasingly, species distribution models based on individual occurrence records. Expert maps are useful at coarse resolution, where they are suitable for delineating unoccupied regions. In contrast, point records typically provide finer-scale occurrence information that can be characterized for its environmental association, but usually suffers from observer biases and does not representatively or fully address the geographical or environmental range occupied by a species.

Innovation We develop a new modelling methodology to combine the complementary informative attributes of both data types to improve fine-scale, large-extent predictions. We use expert delineations to constrain predictions of a species distribution model parameterized with incidental point occurrence records. We introduce a maximum entropy approach for combining the two data types and generalize it to Poisson point process models. We illustrate critical decision making during model construction using two detailed case studies and illustrate features more generally with applications to species with vastly different range and data attributes. Our methods are illustrated in the Supporting Information and with a new R package, bossMaps, that integrates with existing generalized linear modelling and Maxent software.

Main conclusions Our modelling strategy flexibly accommodates expert maps with different levels of bias and precision. The approach can also be useful with other coarse sources of spatially explicit information, including habitat associations, elevational bands or vegetation types. The flexible nature of this methodological innovation can support improved characterization of species distributions for a variety of applications and is being implemented as a standard element underpinning integrative species distribution predictions in the Map of Life (https://mol.org/).

*Correspondence: Cory Merow, Ecology and Evolutionary Biology, Yale University, 165 Prospect Street, New Haven, CT 06520-8106, USA.
E-mail: cory.merow@gmail.com

Keywords
Ecological niche model, maximum entropy, Poisson point process, species distribution model.

INTRODUCTION

Knowledge about species' ranges allows biologists to address ecological and evolutionary problems across large spatial extents, including understanding fundamental processes regarding the spatial variation of different biodiversity attributes and community composition. Similarly, global change scientists, decision-makers and managers require information on species distribution if they are to anticipate patterns and

DOI: 10.1111/geb.12539
http://wileyonlinelibrary.com/journal/geb

consequences of species range changes due to global change (Dawson *et al.*, 2011) or to develop conservation and land management plans. Finally, species distributions underpin several of the targets of the Convention on Biological Diversity to improve the state of biodiversity by 2020 (https://www.cbd.int/sp/targets/), and are considered an essential biodiversity variable that should be carefully quantified and monitored (Pereira *et al.*, 2013). Despite these manifold needs, our understanding of the geographical distributions of most species remains relatively limited, and even for well-studied taxa it remains orders of magnitude coarser than the available environmental data (Jetz *et al.*, 2012a). When georeferenced occurrence data are available they often contain considerable bias in sampling effort (Whittaker *et al.*, 2005; Menke *et al.*, 2009; Boakes *et al.*, 2010; Meyer *et al.*, 2015). Species distribution models (SDMs) are valuable for inferring suitable habitat at unsampled locations, but bias in occurrence records can limit their reliability. For thousands of species, expert range maps have been developed, but they remain too coarse and uneven in their reliable spatial resolution to support most of the requirements of spatial biodiversity science (Hurlbert & Jetz, 2007; Jetz *et al.*, 2008). This has hitherto prevented the development of more reliable, spatially detailed and broadly useful estimates of species true geographical distributions and has imposed strong limits on many spatial biodiversity applications. This 'Wallacean shortfall', or paucity of range information, represents a particularly strong limitation in conservation and spatial planning, where both spatial detail and sufficiently representative sets of species are often key (Lomolino, 2004; Whittaker *et al.*, 2005; Jetz *et al.*, 2012a). In this study, we develop and demonstrate innovations that will allow a significant step toward addressing this gap.

Different data types

We follow up on previous calls for the combination of different sources of biogeographical information, which may allow one to balance their respective strengths and weaknesses and enable the development of improved range predictions (Jetz *et al.*, 2012a; Domisch *et al.*, 2016). Point-level presences are by far the most commonly used data source, as these are readily available in a number of open-access online repositories. Progress has already been made in integrating presence data with other data sources: Fithian *et al.* (2014) and Dorazio (2014) take related approaches to combine presences with presence/absence data for many species to overcome issues of sampling bias; Pagel *et al.* (2014) use presences to bridge gaps in abundance data; and Pagel & Schurr (2012) use presences with demographic data to predict spatial population dynamics. Here, we focus on integrating point-level presences with expert drawn occurrence polygons ('expert maps') to improve distribution models. Our approach can also be extended to expert knowledge on species habitat, elevation or bioclimatic preferences commonly available in field guides or taxon-specific monographs.

For many species, rough maps drawn by single experts or expert groups provide helpful guidelines on species ranges. Such maps may be based on different sorts of published or unpublished occurrence data combined with the experts' own assumptions about environmental and geographical distribution limits. Experts then typically use administrative or rough physiographic or environmental contours to demarcate boundaries. Although valuable across large spatial extents and coarse spatial resolutions (Hurlbert & Jetz, 2007), expert maps have a number of weaknesses. Often hand-drawn or over-generous in extent, they tend to be contiguous 'blobs' that predict false presence at much smaller resolution due to ecologically relevant landscape variation (Hurlbert & White, 2005). This results in a limited spatial accuracy, estimated to only be in the 100–200 km range for many (even well-studied) species (Hurlbert & White, 2005; Hurlbert & Jetz, 2007), that is non-random and varies geographically and ecologically (Jetz *et al.*, 2008). In other words, the species may be absent for large regions inside the expert range maps and present well outside it. For rare or cryptic species, range maps might additionally miss new, disjunct parts of the distribution that have not yet been detected, which is particularly critical. Because they are usually updated with new information only at long and irregular intervals (5–20 year time-scale) expert maps are also prone to quickly become out of date, especially in an age of rapidly growing citizen science data and the new identification of cryptic species/evolutionarily significant units using molecular tools. Finally, due to their very nature, representing a construct of many sources and data types, expert maps also often lack transparency and are not readily reproducible. However, while boundaries may be inaccurately placed, expert maps are usually, overall, an excellent resource for delimiting the broad areas (e.g. >200 km from the boundary) outside which a species is not expected to occur.

In contrast, higher-resolution 'point' observations of occurrence are a more thoroughly specified type of data based on explicit observations in space and time. These data have grown considerably over the last decade, with a particular increase in so-called 'point occurrence' (PO) data (also called 'presence-only' data, as opposed to 'presence–absence' data that typically arise from a designed survey). PO data have low rates of false presence (only arising from errors in identification) compared with expert maps but many more implicit 'false negatives'. Occurrence data have been used, in conjunction with environmental covariates, to build species geographical distribution or environmental niche models (SDMs) that predict species ranges (Guisan & Zimmermann, 2000; Guisan & Thuiller, 2005; Anderson, 2012). But such range predictions are in turn limited by the usually highly non-representative nature of point data (Phillips *et al.*, 2009) and the limited ability of existing methods to address this. One particular concern is their level of overprediction into regions that may be environmentally suitable but out of reach or which lack the appropriate biotic conditions (Soberón, 2007).

Complementary data

With data types such as expert maps offering information about broad geographical range limits and others, such as occurrence points, providing a finer-scale sample to estimate potential environmental associations, opportunities for model-based integration exist. A way to improve our understanding and predictions of species ranges lies in combining such complementary data sources. By identifying observed occurrence–environment relationships, SDMs can quantitatively predict environmentally similar locations based on a specific sample and at spatial resolutions limited only by the associated environmental data (typically much finer than expert maps) (Elith & Leathwick, 2009). Expert maps typically only provide categorical predictions, while SDMs are able to address continuous variation in habitat suitability.

Conversely, the more representative and complete distribution information ingrained in expert maps may help overcome the environmental bias inherent in many PO data sets – typically due to unstructured surveys or variability in detection – which leads to misleading range predictions (Hortal *et al.*, 2008). This bias is impossible to detect from presence data alone; absences are needed but are more difficult to obtain systematically (Merow *et al.*, 2016). Expert maps can address this limitation of PO data to some extent; they characterize absences across large extents and are much less likely to exhibit spatial biases due to the cumulative nature of expert experience informing it and the naturally coarser analysis grain applied. Finally, historical contingencies, dispersal limits and biotic factors that all ultimately constrain the spatial realization of the realized niche are not readily captured by SDM-based niche characterizations from PO data alone, and yet are critical for determining range boundaries (Soberón & Nakamura, 2009). In contrast, expert maps can reflect the outcomes of, for example, competition, local extirpation or dispersal limitation to improve predictions of realized distributions.

Data integration

In this study, we use the concept of spatial offsets in regression models to perform the model-based integration of expert map-like species distribution information with that based on presence-only records. Building on Merow *et al.* (2016), we develop and test a framework for including expert-based range limits as spatial offsets in Poisson point process models (PPPMs) and Maxent-based SDMs. Specifically, we address how well PO data can 'update' knowledge contained in expert maps, via the following questions:

1. Can expert maps help to reduce biased extrapolation in occurrence models?

2. Can point occurrence data reduce false presences predicted inside expert polygons?

3. By combining data types, can we reduce prediction uncertainty, particularly for small sample sizes?

In 'Model', we discuss technical aspects of incorporating expert maps into PPPMs/Maxent. In 'Handling potential data biases', we illustrate how predictions are affected by different types of model specification or bias. In 'Generalizing to other species', we explore applications to other species with a variety of expert map and presence data attributes. Our methods are illustrated in the Supporting Information with a new R package (R Core Team, 2015), bossMaps, that integrates with existing generalized linear modelling and Maxent software. We conclude with a discussion of the strengths and weaknesses of our approach and a prospectus on the applications where it will be most valuable.

Macroecology, biogeography, conservation science, global change projections and spatial decision-making all require reliable species occurrence information. We expect that the proposed integration of expert maps with point occurrence data will enable a wider and more rigorous use of species distribution information in biodiversity science by both elevating the quality of range products overall and extending the number of species with sufficient information to be included in spatial analyses. Integration of the sorts of approaches introduced here in spatial biodiversity informatics infrastructure such as Map of Life (https://mol.org/) holds the promise of continued updating and improvement of species range information as field records and expert knowledge about range limits advance.

METHODS

Model

We provide a joint implementation of our approach to integrating expert maps with PO data for Maxent and PPPMs (cf. Warton & Shepherd, 2010; Chakraborty *et al.*, 2011; Aarts *et al.*, 2012; Fithian & Hastie, 2013; Renner & Warton, 2013; Warton *et al.*, 2014; Renner *et al.*, 2015). These modelling approaches are related; the motivation for using expert maps in PO models comes from minimizing relative entropy (as performed in the popular Maxent software package; Phillips *et al.*, 2006; Merow *et al.*, 2016); however, exploiting an equivalence with inhomogeneous PPPMs simplifies model building and fitting (Merow *et al.*, 2016). Maxent and PPPMs are asymptotically equivalent as the resolution of pixel size decreases, but Maxent imposes a specific way of handling a number of modelling decisions based on tools that are common in machine learning (details in Fithian & Hastie, 2013; Renner & Warton, 2013; Renner *et al.*, 2015). Fundamental modelling assumptions are common to both approaches: the presence locations are independent of one another, the intensity of presence records varies spatially and the intensity varies loglinearly with environmental predictors (Renner *et al.*, 2015). To estimate this intensity, both model the ratio of the density of environmental conditions at m presence locations against those at n background locations, irrespective of whether the background locations were sampled (Merow *et al.*, 2013). This ratio(s) is used to predict the relative occurrence rate (ROR) in cells across a landscape (Merow *et al.*, 2013). It is useful to interpret ROR as a multinomial distribution in geographical space, where the probabilities describe which cells are most likely to contain a presence (Merow *et al.*, 2013). ROR sums to unity across all cells in the landscape, so cells with low ROR may still have a high absolute probability of presence but a lower relative probability than other cells in the landscape. This ROR is the so-called *raw output* of the Maxent software.

Let $P^*(x_i)$ describing the ROR in each cell x_i that we aim to predict based on a vector $\mathbf{z}(x_i)$ of environmental covariates there. Predictions then take the following form (Merow *et al.*, 2013)

$$\ln(P^*(x_i))=\ln(Q(x_i))+\mathbf{z}(x_i)\boldsymbol{\lambda}-\ln(C) \tag{1}$$

where $\boldsymbol{\lambda}$ is a vector of fitted coefficients, $Q(x_i)$ represents a prior expectation of the ROR in cell x_i (referred to as a prior in the entropy literature or an offset in the regression literature and hereafter) and C is a constant that ensures normalization of the ROR across the landscape. In our application, $Q(x_i)$ describes the expert map (see below; other possible applications are given in Merow *et al.*, 2016). The interpretation of the offset is also a multinomial distribution in geographical space (the same as the prediction $P^*(x_i)$) and must be specified in terms of a ROR and its values must sum to unity. Because this model is multiplicative (exponentiate both sides of equation 1), a low value of the offset (outside the expert map) probably implies a low value in the predicted distribution unless the environment is extremely favourable ($\mathbf{z}(x_i)\boldsymbol{\lambda}$ is large). One can think of the term $\mathbf{z}(x_i)\boldsymbol{\lambda}$ as updating the expert map; the PPPM reallocates probability in the offset to cells that better reflect the species' distribution. The coefficients of the model ($\boldsymbol{\lambda}$) thus describe the environmental covariates that differentiate the expert map from the occurrence data, i.e. the environmental bias in the expert map.

An offset, in this case the expert map, reflects information about the species' distribution that is independent of the occurrence records. In other words, if one asked the expert where s/he would expect the next occurrence record to be found, the offset values represent the probability with which the expert would select each cell, e.g. $P(x_i|Y=1)$ where $Y=1$ denotes presence. One important ramification of this design is that the model coefficients do not directly describe the environmental niche of the species (instead they describe which environments the expert missed) and thus cannot be used to make projections into new areas or environments. This issue is discussed in detail in the Discussion.

Defining the offset

A number of considerations are important when converting an expert range polygon into a quantitative gridded offset. Since most expert maps are binary (denoting presence/absence), the key consideration in defining the offset is what values to assign to cells inside the range versus outside it. Because the offset is constrained to sum to unity, this problem can be framed to ask how much cumulative probability one should put inside the expert-defined range ($P_{\text{in}}=\sum_{x\in\text{range}} Q_x$), and how much to assign outside. Varying these probabilities directly affects the relative 'strength' of the expert map in the resulting predictions and can range from P_{in} approaching 1 (assuming a very accurate expert map) to a 'flat' offset in which P_{in} is simply the ratio of the range area to the modelling domain area (assuming that occurrences are just as likely inside versus outside the range map). One way to assign P_{in} is based on omission rates, i.e. the proportion of observed presences outside the expert map, which we refer to as 'expert accuracy'. If we take R as the omission rate for an expert map with m cells inside the range and n cells outside the range, one can assign cells inside the range a prior value of $(1-R)/m$ and values outside the expert range as R/n. Note that this definition also allows one the option to buffer the expert map (discussed below) to quantify which points fall inside the expert map.

There are a number of options for estimating the expert accuracy. Ideally, estimating the omission rate would be done with an independent data set to avoid using the same presence data to estimate the offset as is used in the PPPM. One could use the omission rate for a larger taxonomic ('target') group from the same expert map source to estimate the expert accuracy, on average, across the entire group. An initial analysis for terrestrial vertebrates suggests that only 82% of presence records fell within expert maps (J. Otegui, pers. comm.), while 88% were accurate for 790 freshwater fish in North America (Domisch *et al.*, 2016). If it is not possible to use a target group, one may consider using the same presence points for estimating both the expert omission rate and the PPPM. Though not preferable (because the same data are used for the offset and the PO locations), it is reasonable because the expert map is based on geographical space while the PPPM is fitted in environmental space. However, a possible consequence of using the same presence points to estimate expert accuracy and fit the PPPM is that one can introduce bias by up/downweighting the role of the offset if the presence points happen to be biased. For example, observers may be more likely to report a species outside its expected distribution because it is unusual or unexpected. Or observers may only look for a species within its expected range if their study objective is something other than delineating range boundaries.

The second consideration in defining the offset is the shape of the transition from 'inside' to 'outside' the expert range (Fig. 1). The simplest approach is a step function which drops abruptly at the expert range boundary. However, the spatial resolution of expert maps is typically much coarser (Hurlbert & Jetz, 2007) than the modelling resolution of a SDM. Thus, in most cases, a smooth decay is more representative of the continuously decreasing probability of finding the species at locations further from the expert range. As a generic option that probably covers the spectrum of flexibility that one might need, we suggest that a five-parameter generalized logistic curve can be used to describe smoothing as a function of distance (x) from expert range edge (Richards, 1959):

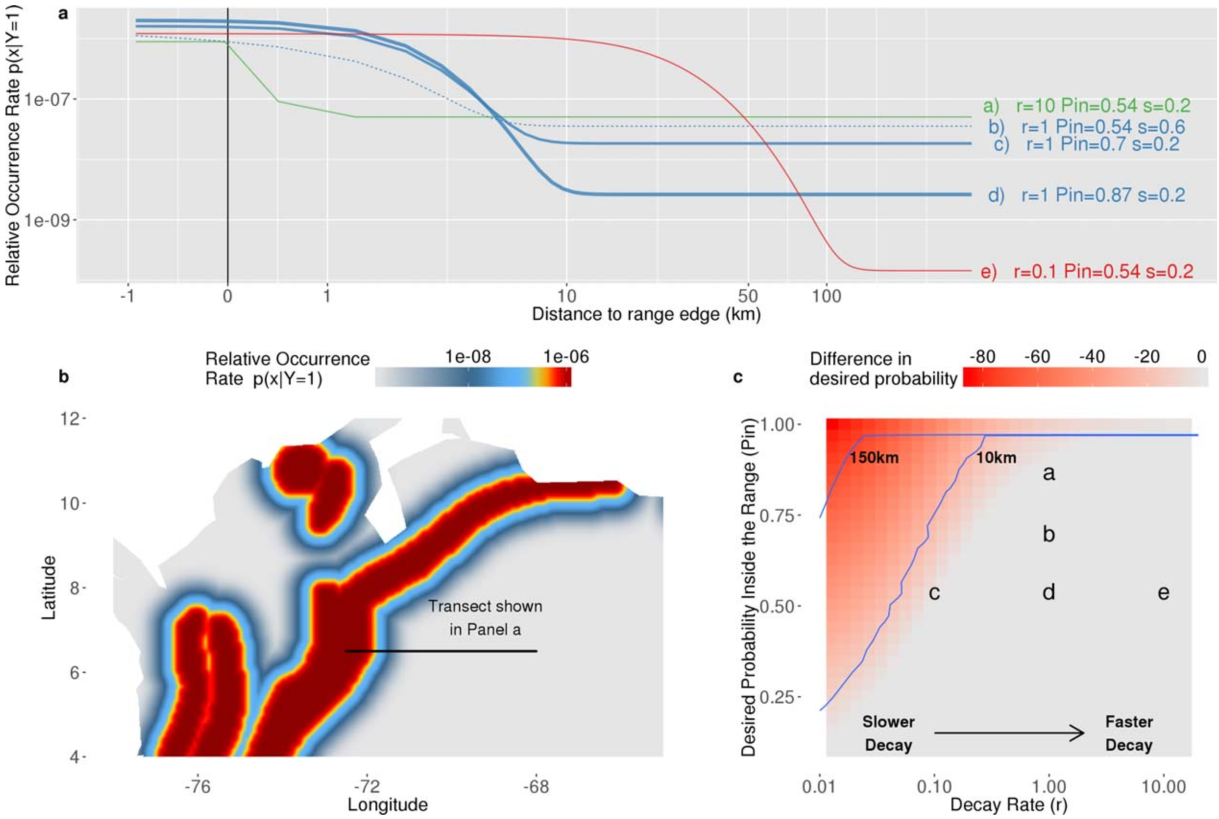

Figure 1 (a) Illustration of various decay curves across the transect shown in panel (b), showing the effects of varying the rate (r), skew (s) and P_{in}. The horizontal axis has an inverse hyperbolic sine transformation and the vertical axis is log-transformed. Letters in labels identify the points shown in panel (c). (b) A subset of the 'expert' range map for the montane woodcreeper (*Lepidocolaptes lacrymiger*) with colours indicating the normalized logistic distance–decay ($P_{\mathrm{in}} = 0.54$, $r = 0.1$, $s = 0.2$) with the upper and lower values estimated to achieve 54% probability inside the expert range. The thick black line shows the transect illustrated in panel (a). (c) Illustration of feasible decay parameters given the range and domain geometry for this species. That is, some parameter combinations (e.g. high P_{in} and slow decay rate r) are mutually exclusive. Colours indicate the difference between the desired P_{in} and the maximum possible P_{in} given the specified curve (deeper reds indicate larger diferences). Blue contour lines show the distance buffer (km) needed around the expert range to include the desired P_{in}. Letters indicate the location of the parameters shown as transects in panel (a).

$$Q(x)=\left(u-\frac{u-l}{\left(1+e^{-r(x-k)}\right)^{1/s}}\right)C^{-1}. \tag{2}$$

The five-parameter logistic is appealing because parameters can be readily interpreted biologically to helps to define useful offsets (Fig. 1). The rate (r) affects the overall decay rate and allows flexibility ranging from a step function (large values) to a flat surface (0). The skew (s) adjusts the symmetry of the decay and ranges from a symmetrical logistic curve to shifting most of the decay outside the expert range (so the probability remains high to the boundary and before decaying). Unlike r, which is rather fundamental for characterizing smoothing, it is reasonable to set $s = 0$ to obtain a simpler model if there is no reason to doubt the expert accuracy inside the map boundary. The shift (k) slides the curve in or out of the expert range to adjust the location but not the shape of the decay. Such a shift can be useful for incorporating a prescribed buffer around the expert map, although in this study we do not adjust this term and effectively consider a four-parameter logistic function. C is a normalization constant ensuring that the ROR sums to unity across the landscape. After specifying the shape of the curve (r, s and k), the final two parameters, the upper (u) and lower (l) asymptotes, are selected via numerical optimization to achieve the desired P_{in} (and relative 'strength' of the prior) for a particular expert range map and modelling domain. These parameters can have considerable effects on predictions (e.g. Figs 2 & 3) and must be estimated based on inferred attributes of the expert maps. That is, they are necessary parameters to adjust to accommodate biologically justified values of r, s and k but are not intrinsically of any biological interest in this application. By using a 'target group' of related species described by the same source of expert maps, it is possible to fit these parameters; however, we anticipate that these parameters will often be set a priori based on an understanding of the expert maps, as in the following examples. While we recommend the five-parameter logistic for its flexibility, we note that a the most minimalistic model to consider would be the three-parameter logistic (setting $s=0, k=0$); further options are considered in the Discussion.

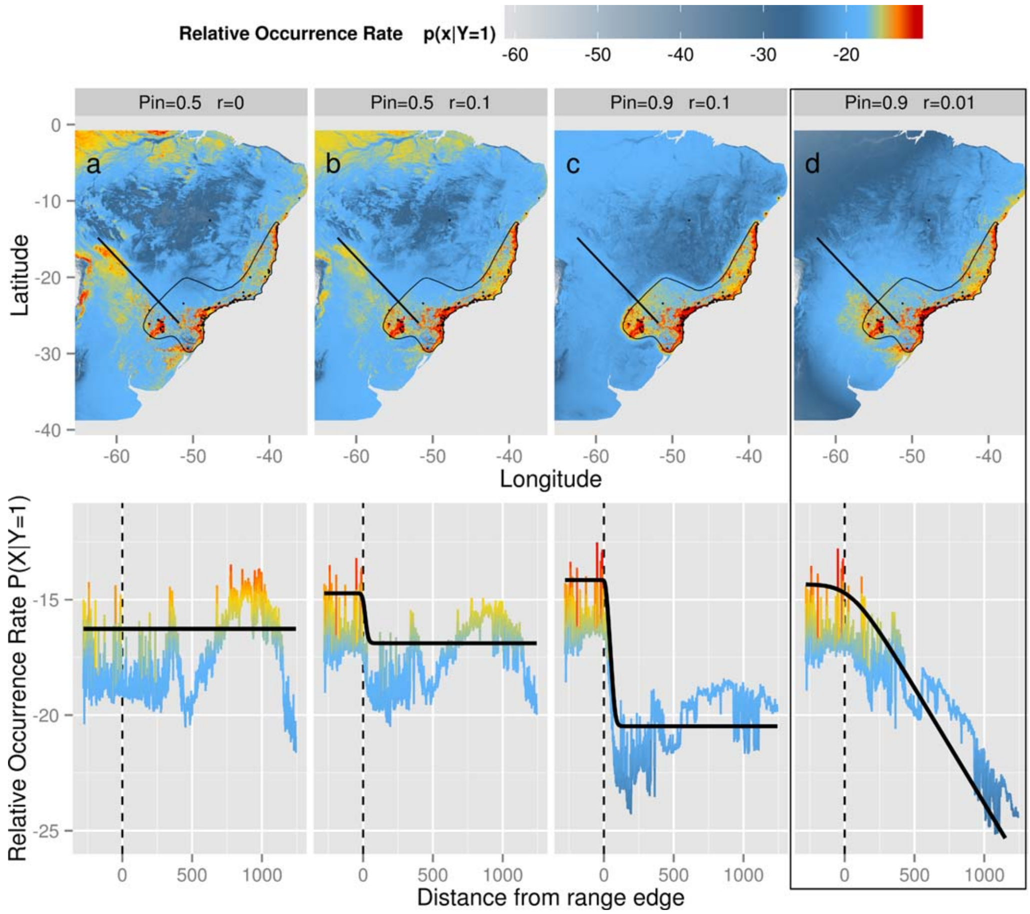

Figure 2 Comparison of various offsets and the resulting predicted relative occurrence rates. Top row: maps of predictions for four offsets ranging from flat (rate, $r = 0$) to a step function ($r = 10$). Bottom row: offset and predictions along the transect (along the dark line in (a)). An illustration of the importance of smoothing the expert maps. If a step function is used (a), equal prior probability is assigned to cells adjacent to the expert boundary and cells many thousands of kilometres away. This assumption is unrealistic because expert errors or dispersal are most likely closer to the expert boundary. By smoothing expert boundaries, we can remove spurious predictions far from the known distribution (based on the critical assumption that the expert maps are already reasonably accurate). Note that the skew parameter s is not varied in this example.

Range and domain geometries impose critical constraints on reasonable combinations of the parameters in equation 2. We denote the smoothing parameter combinations that achieve the desired value of P_{in} as the *feasible set* (grey region in Fig. 1b). That is, certain parameter combinations are not possible because they have mutually exclusive implications for the expert offset. Feasible curves can be evaluated by assessing whether it is possible to achieve the specified accuracy inside the expert range (P_{in}) given a particular decay curve. Figure 1(b) shows allowable combinations of P_{in} and decay rate r for *Lepidocolaptes lacrymiger* in grey. The red scale indicates deviations from the desired value of P_{in} obtained when optimizing the values for the lower and upper asymptotes. The allowable parameter combinations follow intuition – it is paradoxical to select both a high probability inside the expert range (P_{in}) and a very slow decay (r) that puts significant probability outside the range (i.e. upper left corner of Fig. 1b). Conceptually, it is unreasonable to expect that the expert is extremely accurate while at the same time there is high probability outside the range, while numerically it may not possible to assign a large value to (P_{in}) and still distribute the remaining probability over cells outside the expert map without assigning them values so small that they result in numerical underflow.

The combination of parameters in the feasible set depends critically on the modelling domain and spatial resolution. There exists a critical tradeoff between expert accuracy (P_{in}) and spatial decay rate (r): the higher the value of P_{in}, the lower the value of r must be to be part of the feasible set. This derives in part from the normalization constraint that must be imposed on the offset [the normalization constraint derives from the limitation of using presence-only data and predicting ROR (Merow *et al.*, 2013), not from the design of our approach]. For example, if the domain contains an extremely large number of cells outside the expert map (either the domain is very large relative to the expert map or the spatial resolution is very high), it may be impossible to assign high expert accuracy (e.g. 95%) without assigning

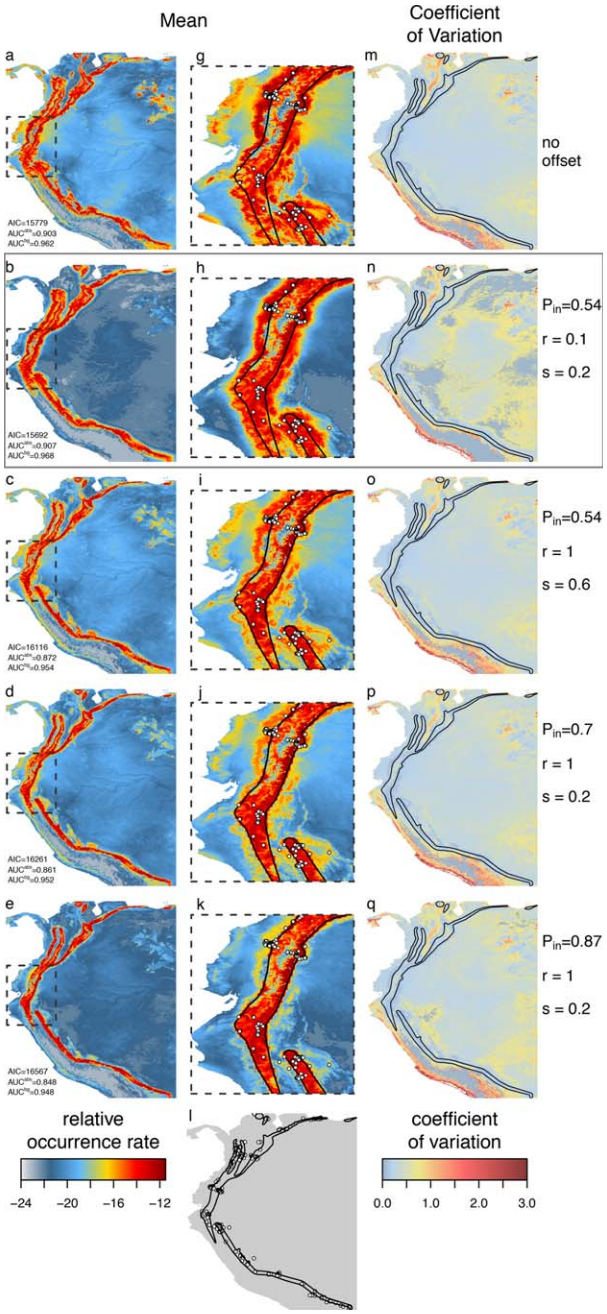

Figure 3 Mapped predictions (left column) and coefficient of variation (right column) to explore the implications of different levels of expert accuracy (P_{in}), rate (r; larger = steeper), and skew (s; larger = softer edge inside the expert boundary). The top row shows a model with no offset for comparison. Rows 2–4 show how predictions change with decay rate r (note $s = 0.6$ is needed when $r = 1$ in row 4 to achieve the desired P_{in}; see Fig. 1b). Rows 5 and 6 show how increasing expert accuracy leads to tighter predictions around the expert range. Notably, a model with an offset in (b) has better AIC and AUC than a model without an offset (a); this model has the 'correct' expert accuracy ($P_{\text{in}} = 0.54$) based on the 374 presence observations available. Notably, reducing the bias by including the expert map does not increase the coefficient of variation, and in some cases, even reduces it. See Fig. D for examples of further parameter combinations.

values outside the expert map values so small that they result in numerical underflow. If the problem derives from a large domain, it is probably appropriate to reduce its size as it is rarely important to employ an expert map while simultaneously inquiring about locations very far from the expert map (the assumption of our modelling method is that the expert map provides useful spatial information). Tips for an appropriate choice of the domain are provided in the Discussion.

The combinations of parameters in the feasible set also depend critically on range geometry. For example, a range with a large perimeter/area ratio (e.g. *L. lacrymiger* in Figs 1 & 3), when spatially smoothed with equation 2, will assign a relatively larger amount of probability immediately outside the expert map, compared with a range with equal area but smaller perimeter/area ratio (e.g. roughly circular). Assigning high expert accuracy to a range with a large perimeter/area ratio may be impossible if the spatial decay rate is slow, because a large amount of cumulative probability would reside immediately outside the expert range, leading to numerical underflow from cells assigned extremely small values far from the expert range. Again, this is a sensible result; one cannot simultaneously suggest that an expert map is highly accurate but that a significant amount of probability resides just outside the expert map. One solution is to simply lower the expected expert accuracy (P_{in}; which must, of course, be justified by data). An alternative is to calculate the cumulative probability as some fixed distance from the expert range boundary; that is, $P_{\text{in}} = 0.95$ might be inappropriate but 95% of the cumulative probability might lie within 100 km of the expert range boundary. By considering locations within this distance as 'inside' the expert range, it is feasible to assign higher P_{in} than would otherwise be possible. As a final solution, if data are available for a large number of species, one could estimate how the decay rate varies as a function of expert map area, perimeter/area ratio or area/domain area ratio.

In summary, it is possible to transform a binary expert range map polygon into a continuous quantitative raster that accounts for the known spatial 'fuzziness' of a typical expert maps at range boundaries. While some parameters are difficult to identify with precision, we note that this process is analogous to 'prior elucidation' in Bayesian approaches, where prior belief is translated into a statistical distribution (Dey & Liu, 2007).

ASSESSING THE METHOD

Data

We provide an intensive assessment for two example bird species, the montane woodcreeper (*L. lacrymiger*, LELA), a bird occupying Andean montane forest, and the solitary tinamou (*Tinamus solitarius*, TISO), a forest species from lowland Brazil. These examples illustrate how predictions are affected by different modelling decisions and attributes of the data. The LELA data consist of: (1) 374 presence observations, (2) 5441 non-detections from the eBird database

(Sullivan *et al.*, 2009), and (3) an expert map based on the International Union for Conservation of Nature (2011) one as used in Jetz & Fine (2012). Point records included in the analysis come from the years 1946 to 2014 and the expert range map is expected to be representative of the *c.* 1980–2010 time frame. The LELA data exhibit some helpful attributes for exploring model construction. Only 54% of presences fall within the expert range, but many are just outside. The range geometry presents some challenges, being relatively small compared with the study domain (the northern portion of South America) and with a high perimeter/area ratio. The LELA data provide a good example of an expert map that can be considerably improved with additional presence data.

The TISO data consist of: (1) 171 presence observations from GBIF (http://gbif.org/) and eBird (Sullivan *et al.*, 2009) accessed through the Map Of Life (https://mol.org) and (2) an expert map from Jetz *et al.* (2012b). The TISO data exhibit some helpful attributes for exploring model construction. Unlike LELA, 92.4% of presences fall within the expert range and the range geometry exhibits a much lower perimeter/area ratio (Fig. 2).

The distribution models were built using 1 km mean annual temperature (MAT; Hijmans *et al.*, 2005) and three satellite-derived cloud-related metrics associated with precipitation (Wilson & Jetz, 2016), mean January and July cloud frequency (CLDJAN and CLDJUL) and cloud seasonality (CLDSEAS).

Model construction

For all tests, we used generalized linear models (glm() in R; R Core Team, 2015) and/or the Maxent software package (Phillips *et al.*, 2006) with settings chosen to fit equivalent models [see the appendices in Renner *et al.* (2015) for options]. Supporting Information Appendix B illustrates our code. We allowed only linear, quadratic and product features/predictor transformations. When fitting models in the Maxent software package, we turned off regularization, removed duplicate presences in the same cell and turned off clamping during model projections. When fitting models with glm() we used downweighted Poisson regression as recommended by Renner *et al.* (2015). All other settings were left at default values. Continuous predictions were assessed with both the presence/absence data and presence/background data using the area under the curve (AUC) and the point biserial correlation. The point biserial correlation showed no qualitative differences from the AUC, so we do not report it. We chose to evaluate against both absence and background data because the absences are inferred from checklist locations very near presence locations and seem more likely to represent non-detections than true absences. Binary predictions were assessed using the true positive rate (TPR; % presences correct), true negative rate (TNR; % absences correct) and correct classification rate (CCR; % presences and absences correctly classified as such). Notably the number of absences (5441) dominates the presences (374) in the calculation of the CCR for LELA. For all tests, we built models both without and with expert offsets with different expert omission rates to examine the effects of different levels of expert accuracy. We also evaluated the predictive performance (based on AUC) of the offset alone to disentangle how improvements derive from the offset versus the presence points.

The effect of offsets

To illustrate how predictions depend on decisions made about the offset, we explored variation in predictions for the different choices of parameters for TISO (Fig. 2) and LELA (Fig. 3). We compared predictions that exclude the expert map with those that incorporate an expert map with different levels of expert accuracy, using all presence data. The TISO models (Fig. 2) focus on varying only P_{in} and r and illustrate an important point about the spatial smoothing enforced by the expert map. Naturally, higher values of P_{in} lead to tighter predictions around the expert range. One might expect that since the expert map only directly contains binary information, an abrupt transition in the offset is appropriate; however, Fig. 2 illustrates the importance of the spatial smoothing that we introduced in equation 2. When an offset is not used, inaccurate predictions for TISO's realized distribution far from the expert map are apparent (Fig. 2a). When an offset is used with an abrupt transition (Fig. 2, last three columns) predictions outside the expert map have relatively lower probabilities; however, many suitable areas are still predicted far from the expert map. This pattern occurs because cells immediately outside the expert map (e.g. 50 km) are assigned the same prior probability as cells 500 or 1000 km from the expert map because the spatial smoothing occurs so abruptly. A more realistic prior expectation is that the probability of observing a presence should continue to decrease with distance from the expert map edge; Fig. 2 (second column) illustrates this smoother spatial decay and it impact on concentrating more probability near the expert map.

For LELA, the presence data suggest that the expert accuracy is 54%; however, to explore the effect of different hypothetical expert accuracy on predictions, we also considered cases where the expert accuracy was assumed to be 70% and 87%. Figure 3 shows predictions for the range of spatial smoothing curves shown in Fig. 1, while Fig. A3 shows a range of other parameter combinations. A number of generalities across the TISO and LELA explorations are apparent:

1. In general, predictions without offsets (Fig. 3, top row) are clearly more diffuse than predictions with offsets. Intuitively, as more confidence is given to the expert map by increasing P_{in}, predictions become more tightly constrained around the expert map.

2. Rate and skew smoothing parameters primarily affected predictions in the immediate vicinity of the expert map (Figs 3, A3), whereas varying P_{in} has a much stronger effect across the domain.

3. When accuracy of the expert map is relatively poor (e.g. LELA; only 54% presences correctly predicted by the expert map) PPPMs with offsets with $P_{\text{in}} = 0.54$ perform best based on the Akaike information criterion (AIC) and AUC (Fig. 3).
4. Including presence data identifies a number of 'holes' in the expert map where habitat is less suitable for this species.
5. A considerable amount of variation in ROR exists inside the expert map.
6. The coefficient of variation is similar in models with and without offsets. Hence reducing bias by including the expert map does not reduce precision. In the case of the large region of low probability in the Amazon for LELA, the absolute value of the prediction standard error is considerably reduced, but since the mean prediction there is also reduced, the coefficient of variation changes minimally. Reducing prediction uncertainty in relatively unsuitable areas is particularly useful for SDM applications.
7. In all cases, the offset alone (i.e. the smoothed expert map) generally predicted far too many absences within the boundaries of the expert map and consequently had much lower AUC (compare Figs 3 & A4). Consequently AUC evaluated against either absences or background points was reduced by *c.* 0.1 compared with models that used presence points. Notably, we do not expect the expert maps to be particularly accurate on their own because their inaccuracies are not due to imprecise boundaries (improved by our smoothing) but rather due to the prediction of false presences.

Handling potential data biases

In order to assess the usefulness of combining expert maps with presence data, we explored scenarios in which different types of bias might exist in either data set. We considered a variety of scenarios by subsampling the full data to explore whether including an offset in PPPMs makes predictions more robust in the face of problematic data. In particular, we wanted to ensure that it was possible to balance the information content of both presences and expert maps such that neither would inappropriately dominate the other if they exhibit considerable bias. We explored extreme cases of biased presences, biased expert maps and the dependence of predictions on presence sample size.

Biased presences

Background. Presence data sets may exhibit considerable bias in records or sampling when they are not the result of systematic surveys (Phillips *et al.*, 2009). Biased samples are particularly problematic for rare or poorly studied species where available records may miss a large geographical and environmental portion of the actual distribution. In such cases, an expert map including these range parts would naturally be useful information for many applications – yet it would preclude benefitting from the information that the small presence point sample may also provide in identifying nearby suitable habitat. Hence, we sought to determine whether bias in presences could persist in PPPM predictions if expert maps were used as offsets.

Methods. We simulated extreme bias by omitting from model fitting presences from large portions of the range to explore whether including expert offsets could improve predictions in portions of the range where the species is not detected. We considered four cases and removed all presences from (1) the southern portion of the range (< 5° S), (2) northern portion of the range (> 5° N), (3) middle portions of the range (< 5° S and < 5° N), and (4) northern and southern portion of the range (< 5° S or > 5° N) from the fitting data. PPPMs were built without offsets and with offsets built with $P_{\text{in}} = 0.54, 0.70, 0.87$ (the same as those shown in Fig. 3) to explore how variation in expected expert accuracy affects predictions (even though we know that the expert accuracy is much lower in this simulation, we expect that one would elect not to use expert maps with only 30% accuracy). AIC and AUC (calculated with either non-detection or background data) were used to compare predictions.

Results. Representative results from omitting presences from the southern portion of the range are shown in Fig. 4; qualitatively similar patterns are seen in other scenarios in Fig. 5. If the expert map is poor (i.e. where only 54% of presences are assumed inside) predictions are more diffuse as in earlier examples. In all cases, predictions were improved based on AIC and AUC when the expert map was given less weight (lower P_{in}). In all cases predictions with $P_{\text{in}} = 0.54$ were superior to models without offsets based on AIC and AUC. Hence, it is apparent that even when presences are completely absent from large portions of the range, the influence of the expert map is sufficient (under a range of offsets) to improve prediction in portions of the range where presences are lacking.

Erroneous expert map

Background. In some understudied systems, expert range maps may also have significant flaws. While we expect that using biased expert maps will have a significant negative impact on predictions, we sought to determine whether PPPMs could still at least correctly highlight areas outside a biased expert map. If models can correctly identify suitable areas outside the expert map, then we would have confidence in both detecting expert map errors and in using offsets in spite of such errors (e.g. with large biodiversity databases, where manual checking of each model is prohibitive).

Methods. To explore the influence of expert map bias, we simulated bias in the LELA expert map by sequentially omitting different portions of the range from model fitting and examining predictions in the omitted portion of the range. More generally, these tests help us to explore how presence points beyond expert boundaries influence predictions. We tested four scenarios: (1) removing the large southern portion of the range; (2) including only the southern portion of the range (i.e. omitting the four northern-most polygons); (3) omitting the large central polygon (leaving only the large

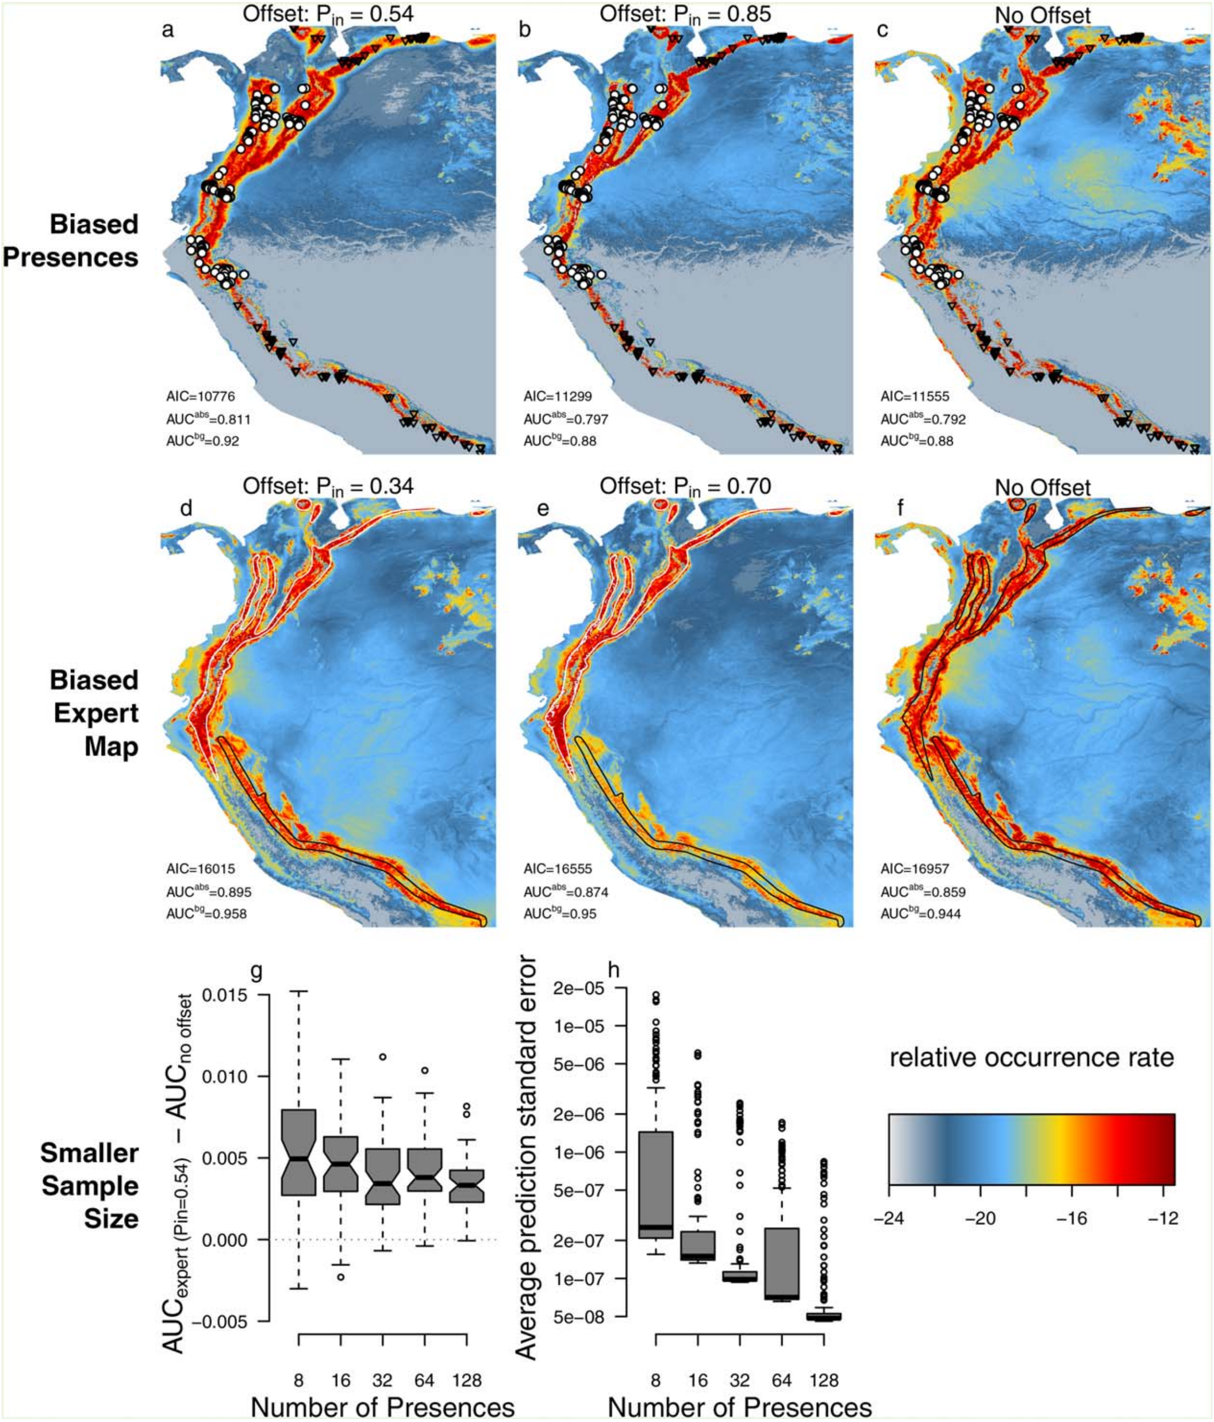

Figure 4 Predictions based on different potential data limitations. Row 1: we simulated bias in the presences by sequentially omitting portions of the presence data from model fitting. Panels (a)–(c) show an example where the southern-most presences (triangles) are omitted from fitting (using only the circles). Suitable habitat was predicted in the south, driven by the expert map offset (see other examples in Fig. A5). Row 2: we simulated bias in the LELA expert map by sequentially omitting portions of the range from model fitting. Panels (d) and (e) show an example where the southern-most portion of the expert range (transparent grey) was omitted from fitting (outlined in black). In spite of expert map bias, suitable habitat was still detected outside the expert map (see other examples in Fig. A6). Row 3: performance comparison of models with and without expert offsets as a function of presence sample size. While diminishing returns were apparent with increasing sample size, it was rare that the expert map did not improve AUC.

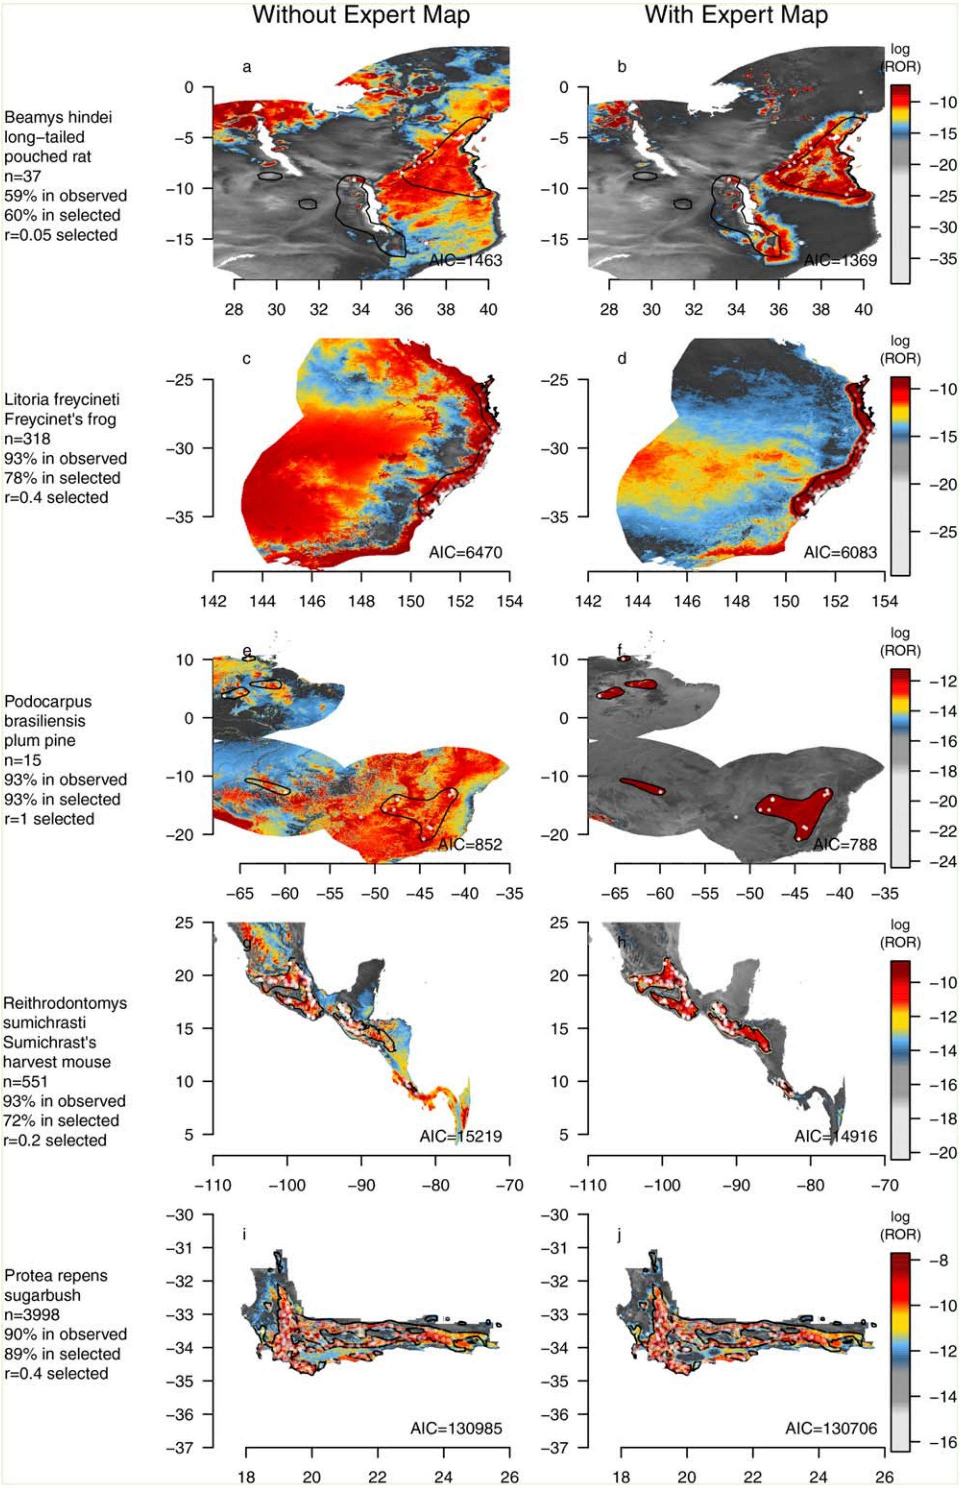

Figure 5 Predictions for species with different range geometries. The left column shows predictions that ignore the expert map (traditional PPPM/Maxent) while the right column shows a model that includes the expert map. The model including the expert map was selected based on having the best AIC among a set of models built with different offsets (the full set of models is shown in Supporting Information Appendix D.2). Beside each row, we indicate the number of presences used for model fitting (n), the percentage of those observed inside the expert map and the values of P_{in} and r used with the expert map. The colour scale is adjusted such that values below the 5% quantile of predictions at presence locations in the model without an expert map are shown on the grey scale while values above this threshold are shown on the colour scale. In all cases, using the expert map improved the AIC.

southern polygons and the three small northern ones); (4) and omitting the three small northern-most polygons. For each case, we compared predictions for different levels of expert accuracy P_{in}; that is, we sought to determine how the ability to predict the true distribution depended upon how much confidence was placed on the expert accuracy. We compared three different offsets for each case, each using the same rate ($r = 0.5$) and skew ($s = 0.2$) parameters, and using $P_{\mathrm{in}} = 0.7$ (relatively high confidence for a poor expert map), $P_{\mathrm{in}} = 0.54$ (the expert accuracy for the full data set, for reference), and the observed value of P_{in} for the portion of the biased expert map used for fitting (0.20–0.34).

Results. In general, it is apparent that even when a large portion of the expert map is omitted from model fitting, the presence points are capable of driving the prediction. We note that this will depend on the number and bias of presence points and the weight assigned to the offset (via P_{in}). Figure 4(d–f) illustrates this pattern for the case where the southern-most expert polygon was omitted from fitting, while Fig. A6 shows the remaining scenarios. In all cases tested, the AIC and AUC were best for models that used the lowest P_{in}, corresponding to the case where the expert map has a smaller influence on predictions. All else being equal (e.g. other smoothing parameters) it is thus preferable to assign a low accuracy to the expert map when it may be biased. However, in spite of the expert bias, the portions of the range which were 'missed' by the expert were still highlighted as relatively more suitable habitat in all cases (i.e. yellow or red in Figs 4 & A6). This observation illustrates that models with offsets can be informative even with imperfect expert maps, so long as undue confidence is not assigned to the expert accuracy.

Sample size

Background. We expected that including an expert map offset would be most helpful when few presence samples were available. Poorly sampled species might be expected to have both poor expert maps and presence samples with considerable bias, hence this is perhaps the most critical scenario for combining data types. We hypothesize that if a large number of presence points are available, the expert map may not add much information unless the presences are spatially biased.

Methods. We randomly subsampled the LELA presence data to obtain 50 replicates of 4, 8, 16, 32, 64 and 128 presences. We built models without offsets and using the expert offsets with $P_{in}=0.54$ and 0.87 used in Fig. 3(b) and (e), respectively. We evaluated all models with AIC and AUC (evaluated with all presence data and either non-detections or background points).

Results. Results were similar across all evaluation measures (Fig. A7), so we focus on the difference in AUC evaluated against non-detection points between models with and without expert offsets (Fig. 4g). Two general patterns are suggestive in Fig. 4(g). Across all sample sizes, models with offsets generally have higher performance than models without ($\Delta\mathrm{AUC} > 0$) Second, the improvement in AUC for models with offsets depends on the number of presences, with greater improvement observed for smaller sample sizes. Similar patterns are observed for AIC, point biserial correlation and AUC based on presence–background data (Fig. A7). Note that the number of presence points required before the AUC asymptotes (around 16 presences here) will vary considerably across species depending on range attributes, sampling biases, etc.

GENERALIZING TO OTHER SPECIES

Background

To better explore the generality of using PPPMs with expert maps we also built models for a set of additional species. For this initial, non-exhaustive exploration, our selection aimed to represent a variety of species groups, regions, unique range geometries and levels of expert accuracy, each posing different challenges for accurate predictions. We sought to determine generally whether including an expert map can typically improve predictions over models that ignore them. Furthermore, we wanted to ensure that parameters describing a range of common patterns of expert accuracy fell within feasible parameter space (cf. Fig. 1c) and that the feasible space was not so small as to preclude useful comparison among models with different spatial smoothing parameters. Finally, we sought to explore whether such potential improvements were readily obtained or required thorough exploration of feasible parameter space. Our goal was to determine if any idiosyncrasies might emerge related to these questions when using expert maps with the following species.

1. *Beamys hindei* (lesser pouched rat) is a rodent with a small, disjunct range with moderate expert accuracy (59%; $n = 37$), wherein many presences are just outside the expert map (http://species.mol.org/species/map/Beamys_hindei).
2. *Litoria freycineti* (wallum rocket frog) is an amphibian species endemic to Australia with a long narrow range with high expert accuracy (93%; $n = 318$) (http://species.mol.org/species/map/Litoria_freycineti).
3. *Podocarpus brasiliensis* is a conifer species from the ancient radiation of podocarps with a small range that is disjunct range across parts of Brazil and Venezuela a large region with apparently high expert accuracy (93%; $n = 15$) (http://species.mol.org/species/map/Podocarpus_brasiliensis).
4. *Reithrodontomys sumichrasti* (Sumichrast's harvest mouse) is a rodent with a relatively large but disjunct distribution that is well characterized by the expert map (93% expert accuracy; $n = 551$) and a large number of presences (http://species.mol.org/species/map/Reithrodontomys_sumichrasti).
5. *Protea repens* is relatively widespread overstorey shrub in the South African mediterranean shrublands with a large range within the shrubland biome. Both the expert map (90% accuracy; $n = 3998$) and the extensive occurrence data set are expected to accurately describe the distribution.

Methods

For each species, we built expert map offsets by varying P_{in} at the observed expert accuracy and values slightly above or below (±0.06–0.08 as appropriate for allowable parameter combinations for each species, cf. Fig. 1c). We also varied values of r to explore different feasible spatial decay parameters of 0.1, 1 and 10. Values for the skew, s, were kept at 0.2 for all models. Our goal was to explore whether improvements over models that ignore expert maps were readily obtained rather than to find the best performing model in the feasible set. Note that the full modelling domain and set of models compared is shown for each species in Supporting Information Appendix D.2, while Fig. 5 shows cropped versions to highlight differences between traditional models that omit the expert map and the best-fitting model (based on AIC) using an expert map.

Results

In general, models including expert maps almost always performed better than those without (Fig. 5), with smoother spatial decay typically leading to better AIC values (Supporting Information Appendix D.2). Importantly, while the offset has a considerable influence on predictions, it does not simply dominate them; unsuitable regions inside the expert map are readily identified by the covariates in all examples and suitable areas outside the expert range are often identified (Fig. 5). Without exhaustively exploring the feasible set of smoothing parameters (cf. Fig. 1c), we readily obtained models with a range of spatial smoothing scenarios that were improvements over models without expert maps. This is encouraging; it may be difficult to precisely estimate expert accuracy for many species, but a comparison of models in Supporting Information Appendix D.2 illustrates that even an approximate estimate typically leads to improvements compared with ignoring expert maps. Increasing P_{in} limited the allowable combinations of smoothing parameters, similar to that observed in Fig. 1c (Supporting Information Appendix E); however, this primarily limited the feasible decay rates (r) when P_{in} was high (>0.90). If the expert map is fairly accurate, it does not seem problematic to avoid strong smoothing to the edges of the expert map (as the implication of using the expert map is that we do not expect presences very far from its boundaries).

Including the expert map was generally useful for excluding large patches of habitat that were otherwise predicted as potentially suitable based on environmental conditions, but where no presences have been observed (Fig. 5, right column). Conversely, the maps shown in Fig. 5 are not particularly good at predicting presences further from the expert boundaries (many occur in grey regions of the maps). This reveals a fundamental tradeoff between reducing false absences and false presences simultaneously, and each may be important in studies with different objectives. Notably, Supporting Information Appendix D.2 shows many examples where smoother offsets better predict points far from expert boundaries; while these models do not have the best AIC values among the models we explored, they were typically still better than models that ignore expert maps. Hence, this framework is flexible enough to accommodate a range of smoothing assumptions to minimize Type I or II errors compared with models that ignore expert maps.

DISCUSSION

Overview

Expert maps and presence-only (PO) data have complementary strengths and weaknesses that make it valuable to combine them in PPPMs and Maxent models. Typically, expert maps have relatively coarse spatial resolution, blurring local habitat variation, and are most informative about where a species is unlikely to be found (Jetz *et al.*, 2012a). In contrast, PO data can help to identify variation in occurrence patterns at high spatial resolution but are less useful in identifying areas where the species does not occur, due to an inability to distinguish absences from non-detection/non-searched areas. Expert maps are typically very good at defining absences beyond range boundaries, which complements PO data. By combining expert maps and PO data, we have shown that: (1) expert maps can improve occurrence predictions when PO data are sparse or biased; (2) presence data can refine expert maps and identify the 'holes' of unsuitable habitat within the range; (3) presence data can identify variation in occurrence patterns in otherwise binary expert maps; (4) improvements from including expert maps are apparent for a range of assumptions about how to characterize the expert map. Hence, combining expert maps with presence data takes advantage of the strengths of different data types to improve predictions of species distributions in a statistically straightforward framework using intuitive, biologically motivated specification of offsets.

Of course, the use of expert maps comes with some caveats that users must recognize. For example, expert maps typically have areas of unsuitable habitat within the range boundaries. These are assigned the same relative probability as suitable areas in the offset, so bias in expert maps carries forward to all subsequent predictions if these patterns are not overcome by sufficient presence data (e.g. Fig. 5b shows considerable variation within the expert map, driven by presences, whereas Fig. 5f does not). Another challenge with using expert maps is determining the appropriate a priori accuracy (P_{in}) to assign to an expert map. This accuracy can be fitted directly from data from a target group of similar species to evaluate the expert's typical accuracy (see 'Outlook and extensions' below); however, the solution for studying a handful of species is less straightforward. If large presence samples are available, expert accuracy can again be estimated from data, with the caveat that any bias in the presence sample will be translated to the expert map.

The predictions provided by our methods should be interpreted as a way to update expert maps. That is, they represent a realized distribution, describing locations that are likely to be currently occupied. While the spatial smoothing imposed by the expert map is valuable for predicting a realized distribution, it is not appropriate for predicting a potential distribution. Spatial smoothing constrains predictions to be 'nearby' the expert map boundary, hence limiting the ability to predict disjunct portions of suitable habitat far from expert boundaries. Consequently, the coefficients of the fitted model do not describe the species' environmental niche (as is commonly inferred in occurrence models) but rather how the species' niche differs from the expert map. This disconnect between geographical distribution and environmental niche means that this modelling framework incorporating expert maps is not suitable for transfer to other locations or times (Merow *et al.*, 2016). Thus the motivation for applying this method is in the refinement of existing expert range maps using point observations.

R package

We have developed an R package (R Core Team, 2015), bossMaps, to facilitate building offsets from expert maps.

The package features tools to explore the shapes of different assumptions about spatial decay (e.g. Fig. 1b) and understand the feasible set of smoothing parameters for a given range geometry (e.g. Fig. 1c) and functions to compute the offset for a given suite of parameters (e.g. Fig. 1a). A tutorial illustrating the full workflow from downloading data to comparing models is provided in Supporting Information Appendix D. bossMaps can be used with any suitable model-fitting software (glm(), Maxent, other options in the appendix of Renner *et al.*, 2015) and includes efficient code to expedite a number of potentially slow computations. All models discussed here were fitted with glm(), choosing appropriate weights to reflect a PPPM as recommended in the appendix of Renner *et al.* (2015). Finally, a number of functions are included to process model output. Model-fitting tips are included in the package documentation, as well as Supporting Information Appendix C, including choosing a domain, exploring the feasible set of smoothing parameters and troubleshooting fitting smoothing parameters for different data scenarios.

Outlook and extensions

A useful option for developing offsets, which alleviates many of the challenges discussed above related to choosing parameters for the offset, is to estimate those parameters based entirely on data. When expert maps come from a common source, one can use the presence data from a large number of species to fit the relative likelihood of finding suitable habitat as a function of distance from the expert range. We expect that there would a general pattern to this relationship that could then be used to estimate a decay curve (described by rate, skew and shift parameters) that could then be used for all the species in that group (or perhaps account for variation in these parameters based on taxonomic or functional subgroups if appropriate). Statistically fitting smoothing functions also provides the opportunity for formal model comparison of different smoothing functions, e.g. three- or four-parameter logistic functions, a Gompertz function or any other sigmoid functions. Exploring other curve shapes that best capture the decay in expert accuracy may be useful, but as the smoothing model is unlikely to be of any intrinsic interest in most studies, we do not view model complexity as a major problem. Hence, we recommend that the flexibility offered by the 5-parameter logistic, with the clear interpretation of r, s, and k (note that set $k = 0$ in this study), appears sufficiently general and suitable to match the amount of information one might expect to extract from a typically coarse resolution expert map.

Extending this framework to use a complete Bayesian formulation would also offer some advantages. In particular, one could incorporate uncertainty in the offset into predictions. Offset uncertainty could be useful to characterize (1) uncertainty in the accuracy of the expert map, which increases in importance for smaller sample sizes, and (2) uncertainty in an appropriate smoothing rate, which is important when expert accuracy may vary by species. To incorporate uncertainty in the offset, one must simply sample from the distribution of values for the offset in each cell, though this may be computationally demanding because it requires estimating a large number of offsets. It is advisable to sample values for all cells simultaneously to ensure that the offset is normalized.

Biased sampling can result from variable search effort, detection probability or both and is particularly worrisome when the sample size is relatively small. The incorporation of expert maps using the approach outlined in this paper is most useful in this situation. Accounting for sampling bias is conceptually straightforward and common in Poisson regression to account for variation in exposure; for example, twice as many presences are expected in environments that have received twice as much sampling effort. This effectively converts a model of intensity (counts) to a model for the *rate* per unit sampling effort. The model can then project the expected intensity for a given (constant) sampling effort. Merow *et al.* (2016) provide worked examples of how sampling bias can be incorporated with other offsets.

While we have focused on expert maps as important sources of spatial information, other similarly coarse spatial models could also be used. For example, known habitat limitations (e.g. elevation, vegetation, ecoregion or land cover) could serve a similar role to expert maps. This information is commonly available in field guides or other species descriptions based on expert knowledge. For example, a species known to occur primarily in forests is probably best interpreted to only rarely be found outside forests (but more common in locations near forests). A spatially smoothed map of forest cover may reasonably describe this distribution in ways that a sparse presence sample cannot. To be fair, we note that objectively choosing the spatial smoothing parameters in such cases may be challenging. Of course, it is preferable to include such information more directly as model covariates *if* sufficient presence data are available to fit the associated coefficients. When presence samples are sparse, however, it may be advisable to incorporate expert information on such habitat characteristics via the offset, so that the presence data can rather inform climatic (or edaphic, etc.) relationships. Furthermore, combining multiple sources of spatial information in a model simultaneously is as simple as multiplying the offsets together and normalizing the result (details in Merow *et al.*, 2016).

Conclusions

Combining expert maps and point occurrence data in models that predict species distributions offers a number of advantages over models that rely on just one data type. Point occurrences, annotated with environmental characteristics, can identify unsuitable locations within the expert-delineated range, and compared with expert maps provide significantly enhanced spatial detail. In turn, the inclusion of expert maps can significantly improve predictions of species distributions based on occurrence points due to their ability to help overcome (1) small or (2) biased presence samples and (3) their ability to account for spatial factors (historical distribution constraints, dispersal limits, biotic factors) not readily described by available covariates. Our approach to using expert maps offers the flexibility to reflect different assumptions (accuracy, smoothing) about expert maps, which can often improve range predictions even when the assumptions are imperfect. Future extension of the method is likely to enable an extension to other types of species distribution data (*sensu* Jetz *et al.*, 2012a), such as expert habitat or elevation restrictions, in a similar statistical framework. While the magnitude of improvements in range prediction will vary and advances gained will often be incremental, the presented approach lends itself ideally to an iterative process where both expert knowledge and occurrence data are updated periodically or even continuously. For example, we envision a scaled-up implementation of the method in the spatial biodiversity infrastructure Map of Life with the potential for ongoing updates to expert-defined limits informed by new data advances and iteratively advanced integrated estimates of distribution. This combination of integrative modelling approaches combining data types, ongoing data flows and infrastructure bringing together data and expert curation charts the path toward a more general and continued advance toward the best-possible knowledge of species distribution for research and application.

ACKNOWLEDGEMENTS

C.M. acknowledges funding from NSF Grants 1565046, 1046328, and 1137366. A.W. acknowledges funding from NASA Biodiversity Grant NNX16AQ45G. W.J. acknowledges funding from NASA Biodiversity Grant NNX11AP72G and NSF grants DBI 0960550, DBI 1262600, and DEB 1026764.

REFERENCES

Aarts, G., Fieberg, J. & Matthiopoulos, J. (2012) Comparative interpretation of count, presence–absence and point methods for species distribution models. *Methods in Ecology and Evolution*, **3**, 177–187.

Anderson, R.P. (2012) Harnessing the world's biodiversity data: promise and peril in ecological niche modeling of species distributions. *Annals of the New York Academy of Sciences*, **1260**, 66–80.

Boakes, E.H., McGowan, P.J.K., Fuller, R.A., Chang-Qing, D., Clark, N.E., O'Connor, K. & Mace, G.M. (2010) Distorted views of biodiversity: spatial and temporal bias in species occurrence data. *PLoS Biology*, **8**, e1000385.

Chakraborty, A., Gelfand, A., Wilson, A., Latimer, A.M. & Silander, J.A. (2011) Point pattern modelling for degraded presence only data over large regions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **60**, 757–776.

Dawson, T.P., Jackson, S., House, J., Prentice, I.C. & Mace, G.M. (2011) Beyond predictions: biodiversity conservation in a changing climate. *Science*, **332**, 53–58 (erratum in *Science*, **332**, 664).

Dey, D.K. & Liu, J. (2007) A quantitative study of quantile based direct prior elicitation from expert opinion. *Bayesian Analysis*, **2**, 137–166.

Domisch, S., Wilson, A.M. & Jetz, W. (2016) Model-based integration of observed and expert-based information for assessing the geographic and environmental distribution of freshwater species. *Ecography*, **39**, 1–11.

Dorazio, R. (2014) Accounting for imperfect detection and survey bias in statistical analysis of presence-only data. *Global Ecology and Biogeography*, **23**, 1472–1484.

Elith, J. & Leathwick, J.R. (2009) Species distribution models: ecological explanation and prediction across space and time. *Annual Review of Ecology, Evolution, and Systematics*, **40**, 677–697.

Fithian, W. & Hastie, T. (2013) Finite-sample equivalence of several statistical models for presence-only data. *Annals of Applied Statistics*, **7**, 1917–1939.

Fithian, W., Elith, J., Hastie, T. & Keith, D.A. (2014) Bias correction in species distribution models: pooling survey and collection data for multiple species. *Methods in Ecology and Evolution*, **6**, 424–438.

Guisan, A. & Thuiller, W. (2005) Predicting species distribution: offering more than simple habitat models. *Ecology Letters*, **8**, 993–1009.

Guisan, A. & Zimmermann, N. (2000) Predictive habitat distribution models in ecology. *Ecological Modelling*, **135**, 147–186.

Halvorsen, R. (2012) A strict maximum likelihood explanation of MaxEnt, and some implications for distribution modelling. *Sommerfeltia*, **36**, 1–132.

Hijmans, R.J., Cameron, S.E., Parra, J.L., Jones, P.G. & Jarvis, A. (2005) Very high resolution interpolated climate surfaces for global land areas. *International Journal of Climatology*, **25**, 1965–1978.

Hortal, J., Valverde, A.J., Gómez, J.F. & Lobo, J.M. (2008) Historical bias in biodiversity inventories affects the observed environmental niche of the species. *Oikos*, **117**, 847–858.

Hurlbert, A.H. & Jetz, W. (2007) Species richness, hotspots, and the scale dependence of range maps in ecology and conservation. *Proceedings of the National Academy of Sciences USA*, **104**, 13384–13389.

Hurlbert, A.H. & White, E.P. (2005) Disparity between range map- and survey-based analyses of species richness: patterns, processes and implications. *Ecology Letters*, **8**, 319–327.

International Union for Conservation of Nature (2011) *IUCN Red List of threatened species*. http://www.iucnredlist.org/

Jetz, W. & Fine, P.V.A. (2012) Global gradients in vertebrate diversity predicted by historical area–productivity dynamics and contemporary environment. *PLoS Biology*, **10**, e1001292.

Jetz, W., Sekercioglu, C.H. & Watson, J.E.M. (2008) Ecological correlates and conservation implications of overestimating species geographic ranges. *Conservation Biology*, **22**, 110–119.

Jetz, W., McPherson, J.M. & Guralnick, R.P. (2012a) Integrating biodiversity distribution knowledge: toward a global map of life. *Trends in Ecology and Evolution*, **27**, 151–159.

Jetz, W., Thomas, G.H., Joy, J.B., Hartmann, K. & Mooers, A.O. (2012b) The global diversity of birds in space and time. *Nature*, **491**, 444–448.

Lomolino, M.V. (2004) Frontiers of biogeography: new directions in the geography of nature. *Conservation biogeography*, (ed. By M.V. Lomolino and L.R. Heaney) pp. 369–419. Sinauer Associates, Sunderland, MA.

Menke, S.B., Holway, D.A., Fisher, R.N. & Jetz, W. (2009) Characterizing and predicting species distributions across environments and scales: Argentine ant occurrences in the eye of the beholder. *Global Ecology and Biogeography*, **18**, 50–63.

Merow, C., Smith, M.J. & Silander, J.A. (2013) A practical guide to MaxEnt for modeling species' distributions: what it does, and why inputs and settings matter. *Ecography*, **36**, 12–11.

Merow, C., Allen, J.M., Aiello-Lammens, M.E. & Silander, J.A. (2016) Improving niche and range estimates with Maxent and point process models by integrating spatially explicit information. *Global Ecology and Biogeography*, **25**, 1022–1036.

Meyer, C., Jetz, W., Guralnick, R. & Kreft, H. (2015) Global priorities for an effective information basis of biodiversity distributions. *Nature Communications*, **6**, 1–8.

Pagel, J. & Schurr, F.M. (2012) Forecasting species ranges by statistical estimation of ecological niches and spatial population dynamics. *Global Ecology and Biogeography*, **21**, 293–304.

Pagel, J., Anderson, B.J., O'hara, R.B., Cramer, W., Fox, R., Jeltsch, F., Roy, D.B., Thomas, C.D. & Schurr, F.M. (2014) Quantifying range-wide variation in population trends from local abundance surveys and widespread opportunistic occurrence records. *Methods in Ecology and Evolution*, **5**, 751–760.

Pereira, H.M., Ferrier, S., Walters, M. *et al.* (2013) Essential biodiversity variables. *Science*, **339**, 277–278.

Phillips, S.J., Anderson, R.P. & Schapire, R.E. (2006) Maximum entropy modeling of species geographic distributions. *Ecological Modelling*, **190**, 231–259.

Phillips, S., Dudik, M., Elith, J., Graham, C., Lehmann, A., Leathwick, J. & Ferrier, S. (2009) Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecological Applications*, **19**, 181–197.

R Core Team. (2015) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.

Renner, I.W. & Warton, D.I. (2013) Equivalence of MAXENT and Poisson point process models for species distribution modeling in ecology. *Biometrics*, **69**, 274–281.

Renner, I.W., Elith, J., Baddeley, A., Fithian, W., Hastie, T., Phillips, S.J., Popovic, G. & Warton, D.I. (2015) Point process models for presence-only analysis. *Methods in Ecology and Evolution*, **6**, 366–379.

Richards, F.J. (1959) A flexible growth function for empirical use. *Journal of Experimental Botany*, **10**, 290–301.

Soberón, J. (2007) Grinnellian and Eltonian niches and geographic distributions of species. *Ecology Letters*, **10**, 1115–1123.

Soberón, J. & Nakamura, M. (2009) Niches and distributional areas: concepts, methods, and assumptions. *Proceedings of the National Academy of Sciences USA*, **106**(Suppl. 2):19644.

Sullivan, B.L., Wood, C.L., Iliff, M.J., Bonney, R.E. & Fink, D. (2009) eBird: a citizen-based bird observation network in the biological sciences. *Biological Conservation*, **142**, 2282–2292.

Warton, D.I. & Shepherd, L.C. (2010) Poisson point process models solve the 'pseudo-absence problem' for presence-only data in ecology. *Annals of Applied Statistics*, **4**, 1383–1402.

Warton, D.I., Shipley, B. & Hastie, T. (2014) CATS regression: a model based approach to studying trait based community assembly. *Methods in Ecology and Evolution*, **6**, 389–398.

Whittaker, R.J., Araújo, M.B. & Jepson, P. (2005) Conservation biogeography: assessment and prospect. *Diversity and Distributions*, **11**, 3–23.

Wilson, A.M. & Jetz, W. (2016) Remotely sensed high-resolution global cloud dynamics for predicting ecosystem and biodiversity distributions. *PLoS Biology*, **14**, e1002415–e1002420.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

Appendix A: Data
Appendix B: The Maximum Entropy Motivation for offsets
Appendix C: Model Fitting Tips
Appendix D: Further Explorations
Appendix E: Understanding Allowable Parameter Combinations
Appendix F: References and Appendices

BIOSKETCH

Cory Merow is a quantitative ecologist interested in forecasting ecological responses to global change.

Editor: Pedro Peres-Neto