

ECOGRAPHY

Research article

Predicting predator–prey interactions in terrestrial endotherms using random forest

John Llewelyn^{1,2}  ^{1,2}, Giovanni Strona^{3,11} , Christopher R. Dickman⁴ , Aaron C. Greenville^{5,6,7}, Glenda M. Wardle⁵, Michael S. Y. Lee^{8,9}, Seamus Doherty^{1,2}, Farzin Shabani¹⁰ , Frédéric Saltré^{1,2}  and Corey J. A. Bradshaw^{1,2} 

¹Global Ecology, Partuyarta Ngadluku Wardli Kuu, College of Science and Engineering, Flinders University, Adelaide, SA, Australia

²ARC Centre of Excellence for Australian Biodiversity and Heritage, Wollongong, NSW, Australia

³Faculty of Biological and Environmental Sciences, University of Helsinki, Helsinki, Finland

⁴Desert Ecology Research Group, School of Life and Environmental Sciences, University of Sydney, Sydney, NSW, Australia

⁵School of Life and Environmental Sciences, University of Sydney, Sydney, NSW, Australia

⁶School of Life Sciences, University of Technology Sydney, Sydney, NSW, Australia

⁷National Environmental Science Program, Threatened Species Recovery Hub, University of Sydney, Sydney, NSW, Australia

⁸College of Science and Engineering, Flinders University, Adelaide, SA, Australia

⁹Earth Sciences Section, South Australian Museum, North Terrace, Adelaide, SA, Australia

¹⁰Department of Biological and Environmental Sciences, College of Arts and Sciences, Qatar University, Doha, Qatar

¹¹European Commission, Joint Research Centre, Directorate D – Sustainable Resources, Ispra, Italy

Correspondence: John Llewelyn (johnslllewelyn@gmail.com)

Ecography

2023: e06619

doi: [10.1111/ecog.06619](https://doi.org/10.1111/ecog.06619)

Subject Editor: Daniel Stouffer

Editor-in-Chief: Miguel Araújo

Accepted 18 May 2023



www.ecography.org

Species interactions play a fundamental role in ecosystems. However, few ecological communities have complete data describing such interactions, which is an obstacle to understanding how ecosystems function and respond to perturbations. Because it is often impractical to collect empirical data for all interactions in a community, various methods have been developed to infer interactions. Machine learning is increasingly being used for making interaction predictions, with random forest being one of the most frequently used of these methods. However, performance of random forest in inferring predator–prey interactions in terrestrial vertebrates and its sensitivity to training data quality remain untested. We examined predator–prey interactions in two diverse, primarily terrestrial vertebrate classes: birds and mammals. Combining data from a global interaction dataset and a specific community (Simpson Desert, Australia), we tested how well random forest predicted predator–prey interactions for mammals and birds using species' ecomorphological and phylogenetic traits. We also tested how variation in training data quality – manipulated by removing records and switching interaction records to non-interactions – affected model performance. We found that random forest could predict predator–prey interactions for birds and mammals using ecomorphological or phylogenetic traits, correctly predicting up to 88 and 67% of interactions and non-interactions in the global and community-specific datasets, respectively. These predictions were accurate even when there were no records in the training data for focal species. In contrast, false non-interactions for focal predators in training data strongly degraded model performance. Our results demonstrate

© 2023 The Authors. Ecography published by John Wiley & Sons Ltd on behalf of Nordic Society Oikos

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

that random forest can identify predator–prey interactions for birds and mammals that have few or no interaction records. Furthermore, our study provides guidance on how to prepare training data to optimise machine learning classifiers for predicting species interactions, which could help ecologists 1) address knowledge gaps and explore network-related questions in data-poor situations, and 2) predict interactions for range-expanding species.

Keywords: ecological network, food web, machine learning, Simpson Desert, species interactions, training data

Introduction

Species interactions are essential for maintaining biodiversity (Thompson 1999) because they determine how energy and nutrients flow among organisms (Traill et al. 2010, Allan et al. 2021) and affect the distribution, survival and abundance of species (Van der Putten et al. 2010). Species interactions also determine how ecological disturbances – such as those involving invasive species, habitat modification, and climate change – reverberate through ecological communities (Tylianakis et al. 2008, Blois et al. 2013). This is because environmental changes that directly alter one or a subset of species in a community can have flow-on effects for other species via interaction networks (Kaneryd et al. 2012, Strona and Bradshaw 2018, Morton et al. 2022). Given the need to predict and mitigate human-driven biodiversity declines, interest in modelling interaction networks for biological communities is increasing (Bohan et al. 2017, Strydom et al. 2021, Strona and Bradshaw 2022). However, most species interactions in most ecological communities have not been documented (Hortal et al. 2015, Jordano 2016, Caron et al. 2022). Indeed, it is usually logistically impossible to measure how all species interact in a community – even when only one type of interaction (e.g. predator–prey) is considered (Brousseau et al. 2018, Strona 2022). In addition to the challenge of identifying long-established interactions, the rapid and widespread translocation of species across the planet due to migratory responses to climate change and movement by humans means we also need to identify interactions that could occur between species that previously did not co-occur, and therefore had not previously had the opportunity to interact (Smith and Phillips 2006, Valdovinos et al. 2018).

To overcome gaps in species interaction data, various methods for inferring interactions using species traits/proxies can be applied, such as quantile-regression trophic niche-space modelling (Gravel et al. 2013, Llewelyn et al. 2022), Bayesian linear modelling (Caron et al. 2022), and several machine learning approaches (e.g. K nearest-neighbour – Desjardins-Proulx et al. 2017, Pichler et al. 2020; neural networks/deep learning methods – Fricke et al. 2022; random forest – Sydenham et al. 2021; and ensemble methods – Poisot 2023). Of these, random forest is a popular method due to its ability to learn nonlinear interactions between predictor variables and predict interactions accurately (Desjardins-Proulx et al. 2017, Laigle et al. 2018, Sydenham et al. 2021). Random forests have been used to predict various interaction types, including plant–pollinator (Ornai and Keasar 2020, Sydenham et al. 2021), parasite–host (Kotula et al. 2021),

and predator–prey interactions (Desjardins-Proulx et al. 2017, Parravicini et al. 2020). They have also been used to identify which species' traits are most important for predicting whether species interact (Dellinger et al. 2019, Klomberg et al. 2022). Using training data consisting of interacting and non-interacting species, random forests can learn how ecomorphological traits, phylogenetic position, and/or taxonomy (as a proxy for phylogenetic position; Laigle et al. 2018) relate to which species interact. However, questions remain regarding the ability of random forests to predict different types of interactions in different taxonomic groups, which species' data/traits are best for inferring interactions, and how to fine-tune training datasets (Poisot 2023). Furthermore, investigating ways to optimise training data for predicting species interactions using random forest will likely help guide procedures for optimising other machine learning methods.

Identifying predator–prey interactions is central to community ecology because these interactions influence both top–down and bottom–up processes such as energy transfer and extinction cascades (Sinclair et al. 2003). Random forests have been applied to predict such interactions, but these applications have largely been limited to fish communities (Parravicini et al. 2020, Strona et al. 2021) and soil invertebrates (Desjardins-Proulx et al. 2017, Laigle et al. 2018). Indeed, the performance of random forests for inferring predator–prey interactions among terrestrial vertebrates has not been assessed despite this group including many species of high conservation concern (Tilman et al. 2017, Cox et al. 2022) and of important ecological function (e.g. ecological engineers, Samson and Knopf 1996).

If we are to model terrestrial ecosystems realistically to predict and minimise the risk of extinction cascades, reliable inference of trophic interactions is essential. However, trophic-related traits and network structure differ among terrestrial vertebrates, fish and invertebrates (Liem 1984, Liem 1990, Shurin et al. 2006, Miller-ter Kuile et al. 2022), suggesting traits and methods (such as random forest and other machine learning techniques) that effectively predict interactions in fish and invertebrates might not necessarily work well for terrestrial vertebrates. These vertebrates are diverse structurally, phylogenetically, and functionally, potentially making the 'rules' determining interactions complex and difficult for machine learning classifiers to identify (whereas just one trait – body size – largely determines trophic interactions in fish due to their gape limitation; Gravel et al. 2013, Morales-Castilla et al. 2015). Furthermore, interaction data for many terrestrial vertebrate species have been recorded

opportunistically or not at all (although detailed diet data do exist for some species; [Pringle and Hutchinson 2020](#)), potentially leading to sparse (i.e. missing most prey for most predators) and taxonomically biased datasets for training and testing random forests. Having sparse interaction records for predators in training data is problematic because machine learning classifiers such as random forests need to be trained with data that include both interacting (e.g. predator and prey) and non-interacting species pairs, with non-interacting species pairs usually assigned by assuming that undocumented interactions imply no interaction (i.e. pseudo-non-interactions generated from presence-only data; [Pichler et al. 2020](#)). When applied to incomplete (sparse) datasets, this process can identify undocumented but truly interacting species pairs as non-interacting pairs (but see [Strona et al. 2021](#)). Importantly, such ‘false non-interactions’ in training data likely reduce the reliability of interactions and non-interactions inferred from random forests and other classifiers. Similarly, random forests and other classifiers trained on datasets that are taxonomically biased might perform poorly when making predictions for species not represented in the training data ([Strydom et al. 2022](#)).

We designed analyses to test whether random forest could accurately predict predator–prey interactions within and between two diverse terrestrial vertebrate classes – birds and mammals – using both a global (GloBI) and a single-community (Simpson Desert, Australia) predator–prey interaction dataset. We also assessed 1) the importance of ecomorphological traits and phylogenetically based latent traits (i.e. as proxies for phylogenetically conserved traits associated with trophic interactions; [Benandi et al. 2022](#)) for inference accuracy, 2) whether prediction accuracy improved by removing predators from the training dataset that had few interaction records (i.e. those with scarce diet records and were, therefore, prone to generating false non-interactions), and 3) the sensitivity of model performance to the quality of the training data. To test sensitivity to the quality of training data, we compared two scenarios that could occur as a result of sparse training data: 1) a reduction in the number of interaction and non-interaction records used to train the models, and 2) a reduction in the number of interaction records and an increase in the number of false non-interaction records used to train the models. We also investigated how these changes affected model performance when they were restricted to different components of the training dataset (i.e. testing the effect of taxonomic bias on performance).

Our results confirm that: 1) random forests can predict interactions in birds and mammals, 2) models trained on interaction datasets including either ecomorphological traits or phylogenetic position (or both) can perform well, 3) removing predators with few recorded interactions from the training data can improve prediction accuracy, 4) models can identify the prey of predators that have no interaction records in the training data, and 5) model performance is especially sensitive to false non-interactions in the training data that involve records of predatory species for which predictions are being made (i.e. focal predators). Our results highlight

the unrealised potential for the application of random forests for predicting interactions among terrestrial vertebrates. Furthermore, they show to which components of the training data random forests are most sensitive, thereby identifying the most efficient ways to increase accuracy via fine-tuning of training data – highlighting steps that would likely also improve species–interaction inference by machine learning classifiers more generally.

Material and methods

We used two interaction datasets combined with ecomorphological and phylogenetic information to test the performance of random forest for predicting predator–prey interactions in birds and mammals (the Supporting information shows flowchart of dataset preparation). We also tested how the performance of these models was influenced by the type of species data used, the filtering of training data by removing predators with few interaction records in those data, and the quality of the training data. The two interaction datasets were: 1) a dataset from the ‘Global Biotic Interactions database’ (globalbioticinteractions.org; [Poelen et al. 2014](#)) supplemented with interaction data from diet studies done in Australia ([Llewelyn 2022](#)), and 2) an ecosystem-specific dataset from the Simpson Desert in Australia that focuses on seven predators for which detailed dietary information is available ([Llewelyn 2022](#)). We focused on birds and mammals rather than on all tetrapods due to the availability of detailed trait ([Wilman et al. 2014](#)) and phylogenetic (vertlife.org) information for birds and mammals (detailed trait databases for reptiles and amphibians are far from complete), thereby providing a wide range of traits for predicting interactions.

Datasets

Global data

We extracted all predator–prey interactions involving birds and mammals from GloBI (accessed 29 May 2020 and 26 June 2020 for birds and mammals as consumers, respectively). However, GloBI contains few records of species found in Oceania (including Australia) relative to other regions of the world: of the 3329 unique species combinations in the records we extracted from GloBI, only 109 involved a predator and prey species both found in Oceania. Supervised learners, including random forests, need to be trained on data that capture the variation present in the species for which predictions are being made, and so geographic and taxonomic biases in training data could reduce prediction accuracy. We therefore added interactions involving Simpson Desert species and their relatives to the GloBI data (Supporting information). We added data for these species by merging two additional datasets with GloBI: 1) predator–prey records from the Simpson Desert that did not involve the seven focal predators (focal predator = species for which we

made predictions), and 2) interactions of the focal Simpson Desert predators with non-Simpson Desert prey elsewhere in Australia (i.e. from diet studies of the focal predators outside the Simpson Desert; no interactions between the focal predators and Simpson Desert species were included in the global/training data). In both cases, we built these datasets by aggregating predator–prey records from diet studies (Supporting information). We then combined these additional interaction datasets with the GloBI data (adding 363 records) to create an enhanced global interaction dataset (Llewelyn 2022).

To predict species interactions using random forest, we required a training dataset that included interacting and non-interacting pairs of species. Like most interaction datasets, our global dataset did not include non-interactions i.e. because relationships are based on presence-only data. We therefore generated pseudo-non-interactions ('non-interactions' henceforth) by taking all possible two-species combinations from the species in our global interaction dataset and removing those combinations documented as 'interacting' (Supporting information). We also removed combinations where either the 'predator' does not consume vertebrates (e.g. strictly herbivorous species could not be potential predators) or the combination involved species not found on the same continent (using the 'rgbif' package in R, www.r-project.org; Chamberlain et al. 2022). We removed these non-interactions because they would not help models predict which sympatric species known predators preyed on (i.e. the type of interactions we aimed to predict in the Simpson Desert). Although the non-interactions we generated would have included many sympatric, truly non-interacting species combinations, they would have also included false non-interactions (species combination assigned as not interacting, but that do indeed interact) as well as allopatric species pairs that would interact if they were sympatric.

Random forests use explanatory variables such as species' traits to learn and predict associations (Desjardins-Proulx et al. 2017). We therefore added trait data to all the predator and prey species in the interaction and non-interaction datasets (Supporting information). These traits included: body mass, time of activity (whether the species displays nocturnal, diurnal, and/or crepuscular activity), ability to fly, resource groups consumed (plants, vertebrates, and/or invertebrates), and percentage use of different habitat strata (eight habitat categories for birds and five for mammals; Wilman et al. 2014, Supporting information). We added phylogenetic information by calculating eigenvectors describing a species' position in the phylogeny (eigenvector mapping; Guénard et al. 2013), a potential proxy for explanatory variables that are phylogenetically conserved. We extracted phylogenies from VertLife.org, and used the MPSEM (Guénard and Legendre 2022) and ape (Paradis and Schliep 2019) R packages (www.r-project.org) to handle phylogenetic data and extract eigenvectors.

We also identified which species pairs in the non-interaction data included a potential prey species from within the predator's preferred prey-size range (Supporting information). We did this so the ratio of non-interaction records from

inside versus outside the preferred size range in training datasets could be adjusted to optimise performance of random forests. This is analogous to optimising species distribution models by adjusting the ratio of pseudo-absences sampled from among versus outside observed occurrences (Barbet-Massin et al. 2012). To identify a preferred size range, we log-transformed (log_e) predator and prey body masses to linearise body–size relationships (Gravel et al. 2013). We then resampled the interaction dataset to create 100 training and testing dataset pairs using all 3692 interaction records in each of the 100 pairs (75:25 interaction records in training:testing datasets).

To each training set we fitted a series of upper and lower quantile regressions capturing the middle 0.99 to 0.8 quantiles at 0.01 quantile increments. We included a prey mass × predator taxonomic class interaction in the regression models to allow different body–size relationships for predatory birds and mammals (Llewelyn et al. 2022). We then combined each of the 100 testing sets with an equal number of species pairs randomly sampled from the non-interaction set (also with body mass log_e-transformed), and applied the quantile regression models (i.e. upper and lower quantile regressions combinations) built in the previous step to the testing sets. Species combinations in the testing sets that included a potential prey species from within the upper and lower quantile ranges derived from the regressions were predicted as prey, and those outside these limits were excluded as prey.

We assessed model performance by comparing predictions to documented interactions (i.e. asking whether the quantile regression models correctly identified which combinations came from the interaction and non-interaction datasets), and used these results to calculate the true skill statistic to identify the optimal quantile regression model for assigning predator–prey interactions. The true skill statistic is a metric that can effectively evaluate the performance of classifiers such as those predicting species interactions (Poisot 2023). A true skill statistic = 1 indicates a model that perfectly predicts interactions and non-interactions, whereas when it is 0, the model performs no better than assigning links randomly. The model capturing the middle 0.84 quantiles performed best (true skill statistic = 0.26), so we used that model to classify species pairs in the non-interaction dataset as those that included a potential prey from within the predator's preferred size range and those that included a potential prey from outside the predator's preferred size range.

Simpson Desert testing data

It is unclear how the accuracy of random forests trained and tested on data including species from many different communities (such as was the case for our global datasets) translates to model performance when applied to single ecological communities (Parravicini et al. 2020). We therefore tested model performance when applied to seven common predators from the Simpson Desert, a community we chose because it has been intensively studied for over 30 years and has detailed information on species composition and the diets of some

predators. We used the species documented at established field sites in the Simpson Desert as the species assemblage (i.e. potential prey), rather than all species documented in the Simpson Desert bioregion, because the former includes sympatric species that have the potential to interact whereas the latter would likely include species that seldom have the opportunity to interact.

The Simpson Desert bird and mammal assemblage from these sites includes 64 vertebrate-consuming bird and mammal species (49 birds and 15 mammals). However, the diets of many of these species have not been studied in detail. When testing the performance of random forests for predicting predator-prey interactions, we therefore restricted the Simpson Desert testing dataset to seven predatory species with detailed dietary information from studies both in the Simpson Desert and other arid regions in Australia. These seven focal predators included four birds – *Aquila audax*, *Falco berigora*, *Hamirostra melanosternon*, *Tyto alba*, and three mammals – *Canis dingo*, *Felis catus* and *Vulpes vulpes*. *Felis catus* (domestic cats) and *Vulpes vulpes* (red fox) are recent introductions to Australia (since European colonisation in 1788), *Canis dingo* (dingo) was brought to Australia by people > 3000 years ago (Fillios et al. 2012), and the remaining focal predators naturally occur in Australia.

To construct the dataset, we aggregated data from published and unpublished diet studies (excluding species consumed as carrion) of the focal predators in the Simpson Desert (Supporting information, Llewelyn 2022). We also included records from diet studies done outside the Simpson Desert involving species pairs found in the Simpson Desert (i.e. we assumed that if trophic interactions occurred between the species in other communities, they would also occur in the Simpson Desert; Llewelyn 2022), and records involving the focal predators and Simpson Desert species in the global interaction dataset (we removed these records from the global dataset). None of the studies we used to build the Simpson Desert interaction dataset reported non-interactions. We therefore generated non-interactions by assuming any species pairs involving the focal predator and the species recorded from the Simpson Desert field sites (157 bird and mammal species) that have not been documented as predator and prey are not predator and prey. We added the same trait and phylogenetic information to the Simpson Desert data as described for the global dataset. However, we did not include a column indicating if non-interacting species pairs in the Simpson Desert data involved potential prey from within the predator's preferred size range because we only used the Simpson Desert dataset to test model performance (i.e. we did not need to optimise the sampling of these non-interactions because they were not used for training random forests).

Modelling

We used the `ranger` package (www.r-project.org, Wright and Ziegler 2017) in R and the Flinders University high performance computer (Flinders University 2021) to apply

random forest using the model-based recursive partitioning algorithm (Breiman 2001). In the first phase, we optimised the random forest models and tested how performance was affected by 1) the number and choice of variables included in the random forest, and 2) removing predators that had few documented interactions in the training data (i.e. those whose diet is likely not adequately captured in the data). In the second phase, we tested how changes to the quality of the training dataset – including removing records and increasing the number of false non-interactions – influenced model performance (Supporting information shows flowchart of modelling steps and for what different models were used). When optimising models, we assessed model performance by comparing predictions to documented interactions and used these results to calculate the true skill statistic (we also calculated other performance metrics to confirm our results; Supporting information). A true skill statistic of 1 indicates a model that perfectly predicts interactions and non-interactions, a score > 0.75 is 'excellent', between 0.4 and 0.75 is 'fair to good', and < 0.4 indicates 'poor' performance (Fleiss et al. 2003).

Model optimisation and variables

We used the global dataset and cross-validation to optimise training datasets and model parameters, and we compared performance of models trained using: 1) only ecomorphological traits, 2) only phylogenetic eigenvectors, or 3) both types of data (Supporting information). We also compared performance of models that incorporated many vs few explanatory variables/traits to test whether more variables resulted in higher accuracy or overfitting. There were 42 variables in the many-variable ecomorphological model (21 each for predator and prey; Llewelyn 2022) because there were 21 ecomorphological variables in the databases we built (Supporting information). To be consistent, we used 42 phylogenetic eigenvectors in the many-variable phylogenetic model (first 21 phylogenetic eigenvectors for predator and prey). In the few-variable models, we included 10 variables (a number we chose arbitrarily that is substantially smaller than the 42 variables in the many-variable models). These 10 variables were the most important ecomorphological or phylogenetic variables (five each for predator and prey) for predicting interactions, as determined by permutation, which measures the impact of changing a variable's value on model accuracy. To identify the most important variables, we sequentially removed the least-important variable and repeated the permutation and removal steps until 10 variables remained (retained variables for each model listed in the Supporting information, Llewelyn 2022). Although a variable's importance in predicting interactions does not imply causation (Arif and MacNeil 2022), knowing which traits are most important for predicting interactions can help researchers avoid spending time and resources collecting data on traits that are uninformative. The many- and few-variable models that incorporated both types of data simply combined the variables from the respective ecomorphological and phylogenetic models (i.e. either 84 or 20 variables in total).

We also tested whether performance was enhanced or eroded by removing records from the global dataset involving predators with few (< 5) interaction records. We posited that predators that had only a small proportion of their diet documented in the dataset could bias results and lead to the generation of many false non-interactions because, for example, models might be unduly influenced by predators that had not been studied in detail – rather than matching of predator and prey traits. Removing records involving these predators reduced the number of interactions from 3692 in the original global dataset to 3424 in the modified global dataset, and reduced the number of non-interactions from 454924 to 89184.

To optimise training datasets and model parameters, we divided the global dataset's interactions (original and modified data) into training and testing datasets (ratio: 75:25 training:testing), and combined them with non-interaction data. The number of non-interaction records added to the testing datasets equalled the number of interaction records in that dataset, whereas we varied the number of non-interactions in the training datasets to identify the optimal (highest true skill statistic) ratio. We replicated the process of generating training and testing datasets each time we applied a random forest in the parameter-optimization and cross-validation steps described below.

We weighted non-interactions such that their combined weight was equal to that of the observed interactions because models built with weighted non-interactions performed slightly better than those built using unweighted non-interactions (Supporting information; similar to results from species distribution modelling; [Barbet-Massin et al. 2012](#)). These weightings determine the probability an interaction/non-interaction is selected in the bootstrap samples for the trees built from the training datasets. We optimised the training datasets and random forests in terms of 1) number of non-interactions in the training data (above), 2) the ratio of non-interactions from inside vs outside preferred size range in the training data, 3) the probability threshold for assigning interactions – because random forests can estimate probabilities for binary responses such as interact vs not interact ([Malley et al. 2012](#)), 4) the number of variables randomly sampled as candidates at each split in the trees, 5) the number of decision trees in the random forest, and 6) the maximum depth of the trees.

We optimised these parameters by iterating through values for each parameter individually, running five random forests at each value (on five resampled training and testing datasets) and calculating the mean true skill statistic for models at each value. We identified the parameter values that gave models with the highest true skill statistic. We separately optimised the 12 models that differed in terms of included variables (i.e. the many- versus few-variable ecomorphological, phylogenetic, and combined data models) and whether the original or modified global dataset was used (the Supporting information shows optimised parameter sets). We then compared performance of these optimised models by applying them to 100 training and testing datasets generated using the global

dataset cross-validation design described above, and calculated the true skill statistic for each of these 100 cases.

After comparing performance of the twelve models with the global dataset, we tested their performance when trained on the global dataset (with hyper-parameters set to optimise performance, as identified in the previous step) and applied to the Simpson Desert data (Supporting information). Again, we trained and applied 100 random forests for each of the twelve models, and assessed their performance with the true skill statistic. We resampled non-interactions for the training dataset in these 100 analyses because there were more non-interactions than needed, whereas we used all available interaction records in the global datasets (3692 and 3424 for the original and modified global datasets, respectively) in each of the 100 training datasets.

Quality of training dataset and model performance

To test the influence of the size of the training dataset, taxonomic coverage, and prevalence of false non-interactions (true interactions classified as non-interactions) on model performance, we reran the analyses with the optimised model that performed best on the Simpson Desert data (i.e. the many ecomorphological and phylogenetic variable model trained on the modified global dataset; Supporting information) but varied the number of records, taxonomic coverage, and prevalence of false non-interactions in the training datasets (Supporting information).

We deliberately reduced the quality of training data in eight ways to test specific sensitivity hypotheses. We modified the entire training dataset by 1) randomly removing records and 2) randomly switching interaction records to non-interactions (false non-interactions). We hypothesised that model performance would be sensitive to removal of training data because having more records helps the models learn associations between species traits and trophic links, and we hypothesised that models would be sensitive to false non-interactions because random forests would learn the wrong association when trained with erroneous data. We also tested what impact data removal and false non-interactions had when they were restricted to subsets of the training data that included either: 3, 4) focal prey, 5, 6) focal predators, or (7, 8) interactions where neither prey nor predators were focal/Simpson Desert species. We hypothesised that model performance would be especially sensitive to changes involving focal predators and prey because these records involve the species for which predictions are made. Here, the associations between these species' traits and trophic interactions should be especially important for the model to learn. Furthermore, we hypothesised that model performance would be sensitive, albeit less so, to removing records and increasing the proportion of false non-interactions in non-Simpson Desert (non-focal) interactions because these interactions teach the random forest trait-trophic associations that also apply to Simpson Desert species. Assessing the sensitivity to changes in the quality of training data involving the focal species is important because it assesses the suitability of random forest

for inferring interactions of species that have few/no interaction records, a common phenomenon.

For each of the eight ways we reduced training dataset quality, we made these changes at 1% increments relative to the initial number of records in the modified group (e.g. 1% increments of the total number of interactions involving the focal predator), generating 10 training datasets at each increment. We trained the random forests on these reduced-quality datasets, applied them to the Simpson Desert data, and calculated true skill statistics 10 times at each increment. We plotted results to make visual comparison of the sensitivity of model performance to each of the eight types of artificial variation in the quality of training data, and we compared model performance when all relevant data were modified to quantify the impacts of each type of modification.

We also examined the mechanisms underlying changes in model performance by testing how removing records or increasing false non-interactions for the focal predator influenced the: 1) relative probability (or suitability) of interacting assigned by the random forest to each potential prey for each predator (i.e. asking whether changes to the training dataset affected which species were identified as the most likely prey), and 2) mean probability assigned to potential prey for each predator (i.e. asking whether changes to the training dataset cause an overall increase or decrease in the probabilities assigned to potential prey). If lower model performance reflects a systematic shift in all the probabilities (upward or downward), rather than a change in which potential prey are identified as the most likely, the models could still be used to predict a predator's prey because the rank order of probability of different prey remains similar irrespective of absolute model performance and absolute probability.

To test for changes in relative prey suitability, we measured the correlation (Pearson's r) between prey probabilities assigned by models trained on reduced-quality training datasets with that assigned by models trained on the full training dataset (Supporting information). To test for changes in mean probability assigned to potential prey overall, we subtracted the mean probability assigned to potential prey for each predator as calculated by the full training-data models from that calculated by the models trained on the reduced-quality training datasets (Supporting information). For both analyses, we reduced the quality of the training dataset at 1% increments (10 replicate training datasets/models at each increment) and plotted the results to compare visually. We hypothesised that removing records and including more false non-interactions in the training dataset would disrupt the assignment of relative suitability and reduce assigned probabilities overall.

Results

Variables and model performance

Random forests trained and tested on the global data had mean true skill statistics ranging from 0.53 to 0.77 (Fig. 1a, c,

e, g), indicating 'good' to 'excellent' performance (Fleiss et al. 2003). These models tended to perform better when we used the original global interaction dataset rather than the modified dataset that had predators with few prey records removed (mean difference in true skill statistic between models = 0.14; Fig. 1a, c versus e, g). Using few or many variables in the phylogenetic or phylogenetic/ecomorphylogical variable models had little effect on performance when applied to the global datasets (difference in mean true skill statistic < 0.02 in all cases; Fig. 1), but there was a slight decrease in performance associated with using few variables in the ecomorphological-only models (difference in mean true skill statistic = 0.08 and 0.06 between the many- and few-variable ecomorphological models when the original or modified global dataset was used, respectively; Fig. 1a versus c, e versus g).

The models trained on the global dataset and applied to the Simpson Desert data yielded mean true skill statistics between 0.25 and 0.45, indicating 'poor' to 'fair' performance (Fleiss et al. 2003; Fig. 1). Those that used both phylogenetic and ecomorphological variables performed well when trained on the modified global dataset, yielding mean true skill statistics of 0.43 ± 0.02 (mean \pm standard deviation) and 0.45 ± 0.02 when we used few or many variables, respectively (Fig. 1f, h). Using few versus many variables did not strongly affect the performance of the phylogenetic-only or phylogenetic- and ecomorphological-variable models (difference in mean true skill statistic < 0.07 in all four cases; Fig. 1). In contrast, the few-variable ecomorphological models did not perform as well as their many-variable counterparts (difference in mean true skill statistic > 0.13 between the many- and few-variable ecomorphological models when using either the original or modified global datasets Fig. 1b versus d, f versus h).

Using the model that performed best on the Simpson Desert data (i.e. the many-variable model with phylogenetic and ecomorphological traits trained on the modified global dataset; Fig. 1f), performance predicting interaction for recently introduced predators – cats and foxes – was better than that for native predators (true skill statistic = 0.56 ± 0.03 and 0.41 ± 0.02 for recently introduced and native predators, respectively; Fig. 2).

Effects of training dataset size, false non-interactions, and taxonomic coverage

We altered the global (modified) training dataset by randomly removing records and changing interaction records to false non-interactions, applying these changes both to the entire training dataset as well as to three subsets: 1) records involving the focal prey (i.e. species found in the Simpson Desert), 2) records involving the focal predators (i.e. the seven Simpson Desert predators for which we predicted interactions), and 3) records involving species not found in the Simpson Desert. Removing records reduced performance when modifying either the entire training dataset (Fig. 3a; true skill statistic reduced to 0 when all records removed because random forests cannot run without training data)

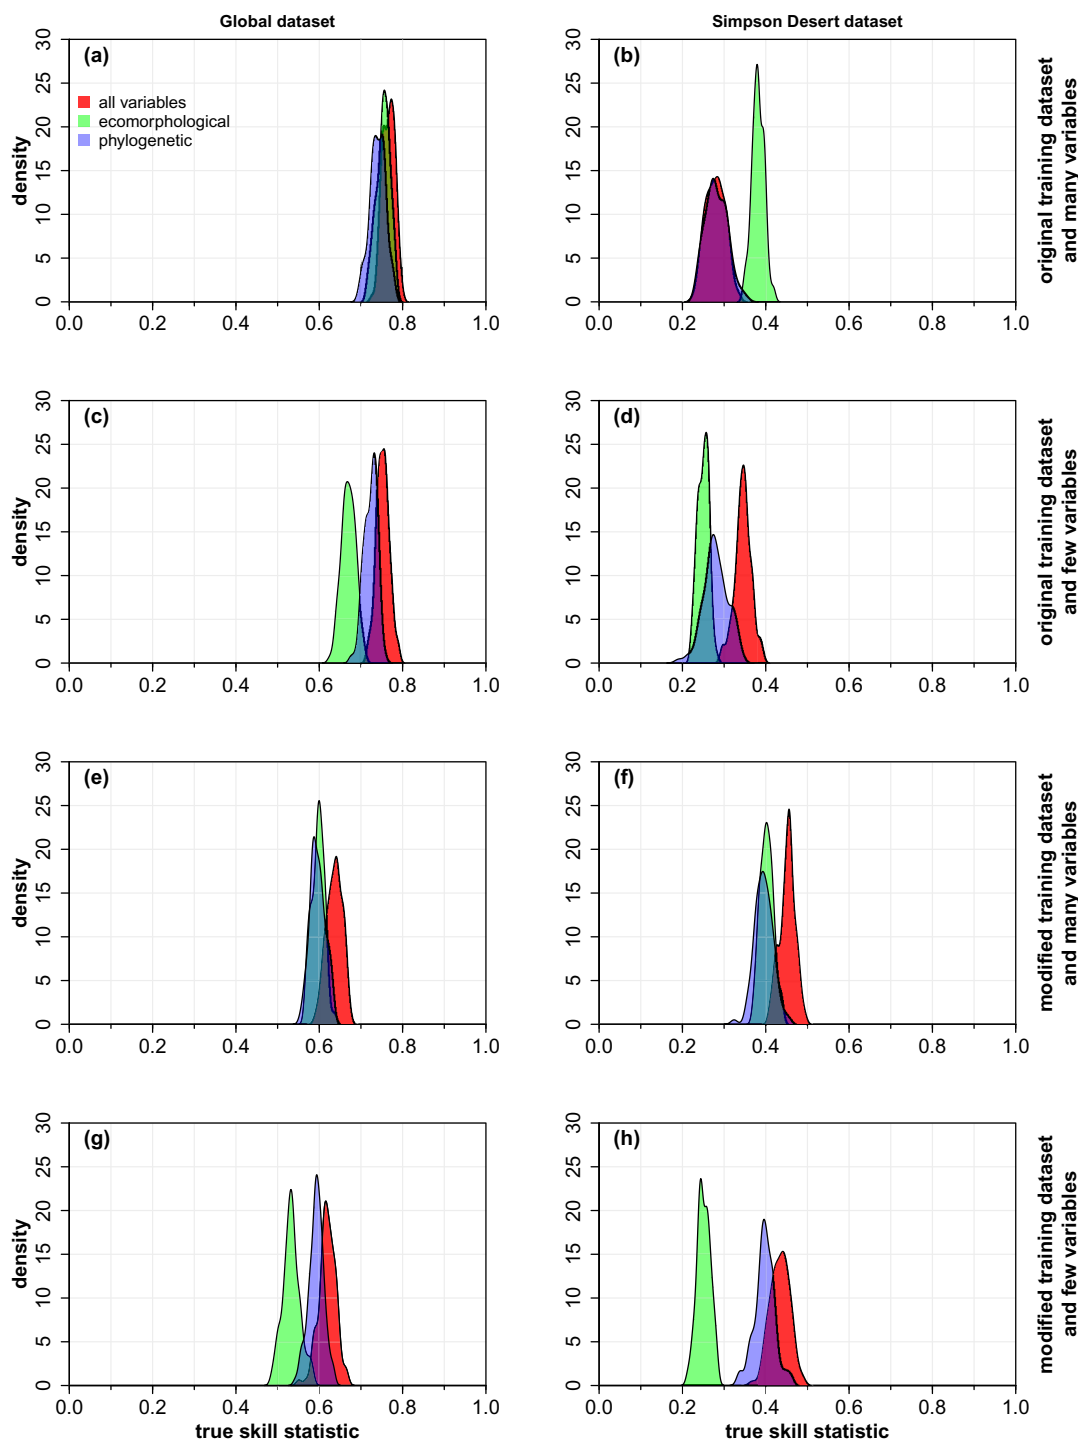


Figure 1. Density plots comparing the performance (true skill statistic) of random forest models predicting predator–prey interactions in birds and mammals. The random forests used either ecomorphological traits, phylogenetic eigenvectors, or both types of data, and were applied to two interaction datasets: a global interaction dataset (panels in the left column) and a specific ecosystem dataset (seven predators and their potential prey/sympatric species in the Simpson Desert; panels in the right column). We also assessed the effects of using many variables (42, 42, and 84 variables, for ecomorphological, phylogenetic eigenvector, and both data-type models) or few variables (10, 10, and 20 variables, for ecomorphological, phylogenetic eigenvector, and both data-type models), and restricting the training dataset to predators with ≥ 5 prey records (modified training dataset). (a) and (b) show results from models that use the complete training dataset and many variables; (c) and (d) show results from models that use the complete training dataset and few variables; (e) and (f) show results from models that use the modified training dataset and many variables; and (g) and (h) show results from models that use the modified training dataset and few variables.

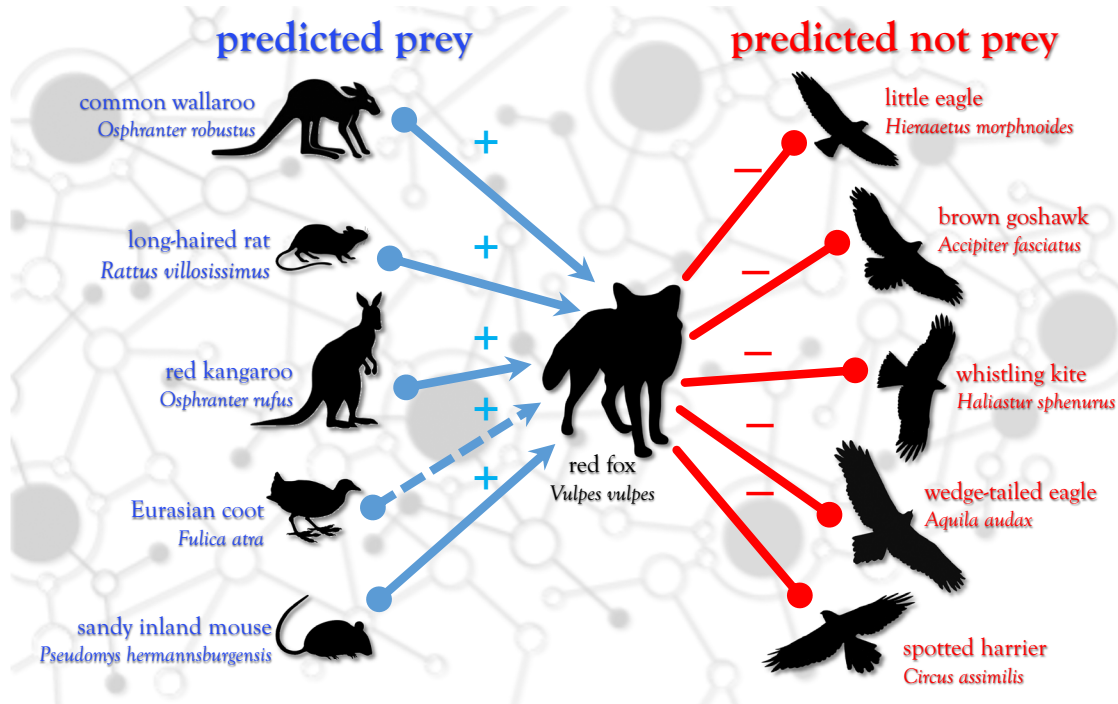


Figure 2. The five most- and least-likely mammalian and bird prey for the red fox *Vulpes vulpes* in the Simpson Desert predicted by random forest. Solid blue arrows indicate prey that were predicted and observed, dashed blue arrows show prey that were predicted but not observed (potentially indicating interactions that occur but have not yet been recorded), and red lines indicate species that were not predicted or observed as prey. The five most-likely prey according to the many-variable phylogenetic and ecomorphological model (run 100 times and trained on the modified training data) were the common wallaroo *Osphranter robustus*, long-haired rat *Rattus villosissimus*, red kangaroo *Osphranter rufus*, Eurasian coot *Fulica atra*, and sandy inland mouse *Pseudomys hermannsburgensis*. The five least-likely prey were the little eagle *Hieraetus morphnoides*, brown goshawk *Accipiter fasciatus*, whistling kite *Haliastur sphenurus*, wedge-tailed eagle *Aquila audax* and spotted harrier *Circus assimilis*.

or the focal-predator component only (Fig. 3e; true skill statistic reduced to 0.24 ± 0.01 [mean \pm standard deviation] when all focal-predator records were removed). However, these effects became pronounced only when $> 80\%$ of the interaction and non-interaction records were removed (Fig. 3a, e), which equates to removing > 15750 or 1240 records from the entire or focal-predator components of the training datasets, respectively. Removing records involving the focal prey or non-Simpson Desert species had only weak or no effect on model performance (Fig. 3c, g; true skill statistic = 0.36 ± 0.02 and 0.46 ± 0.01 when removing all focal prey or all non-Simpson Desert species, respectively).

The effect of false non-interactions on model performance was stronger than that of removing records. False non-interactions assigned at random to interaction records in either the entire training dataset or the component of the training dataset involving focal predators degraded model performance more rapidly (Fig. 3b, f; true skill statistic reduced to 0 and 0.04 ± 0.01 for models when changing all records or all records involving focal predators to non-interactions, respectively). These declines were triggered when changing $> 40\%$ of the interaction records in the relevant group to false non-interactions (Fig. 3b, f), which equates to 1370 records in the entire training-dataset analysis and 235 in the focal-predator analysis. False non-interactions in the training data

involving the focal prey (Fig. 3d) or non-Simpson Desert species (Fig. 3h) caused slight declines in model performance (true skill static reduced to 0.36 ± 0.03 and 0.34 ± 0.02 when changing all interactions records to false non-interactions for focal prey species and non-Simpson Desert species, respectively).

Focal-predator training data and the mechanisms underlying changes in model performance

The correlation between probabilities assigned to potential prey by models trained on the full versus record-removed global datasets showed a gradual decrease as more interactions involving the focal predator were removed from the training data, and this decline became more rapid when removing $> 80\%$ of records (Fig. 4a). The correlation between these probabilities reduced to 0.59 ± 0.02 when all records involving focal predators were removed from the training dataset, indicating that models identified many of the same species as prey even when there were no records involving the focal predator in the training data. The mean probability assigned to potential prey for each predator was reduced by removing records of the focal predator from the training data and, like the correlation results, this change was small until 80% of records were removed (Fig. 4c). After removing all records,

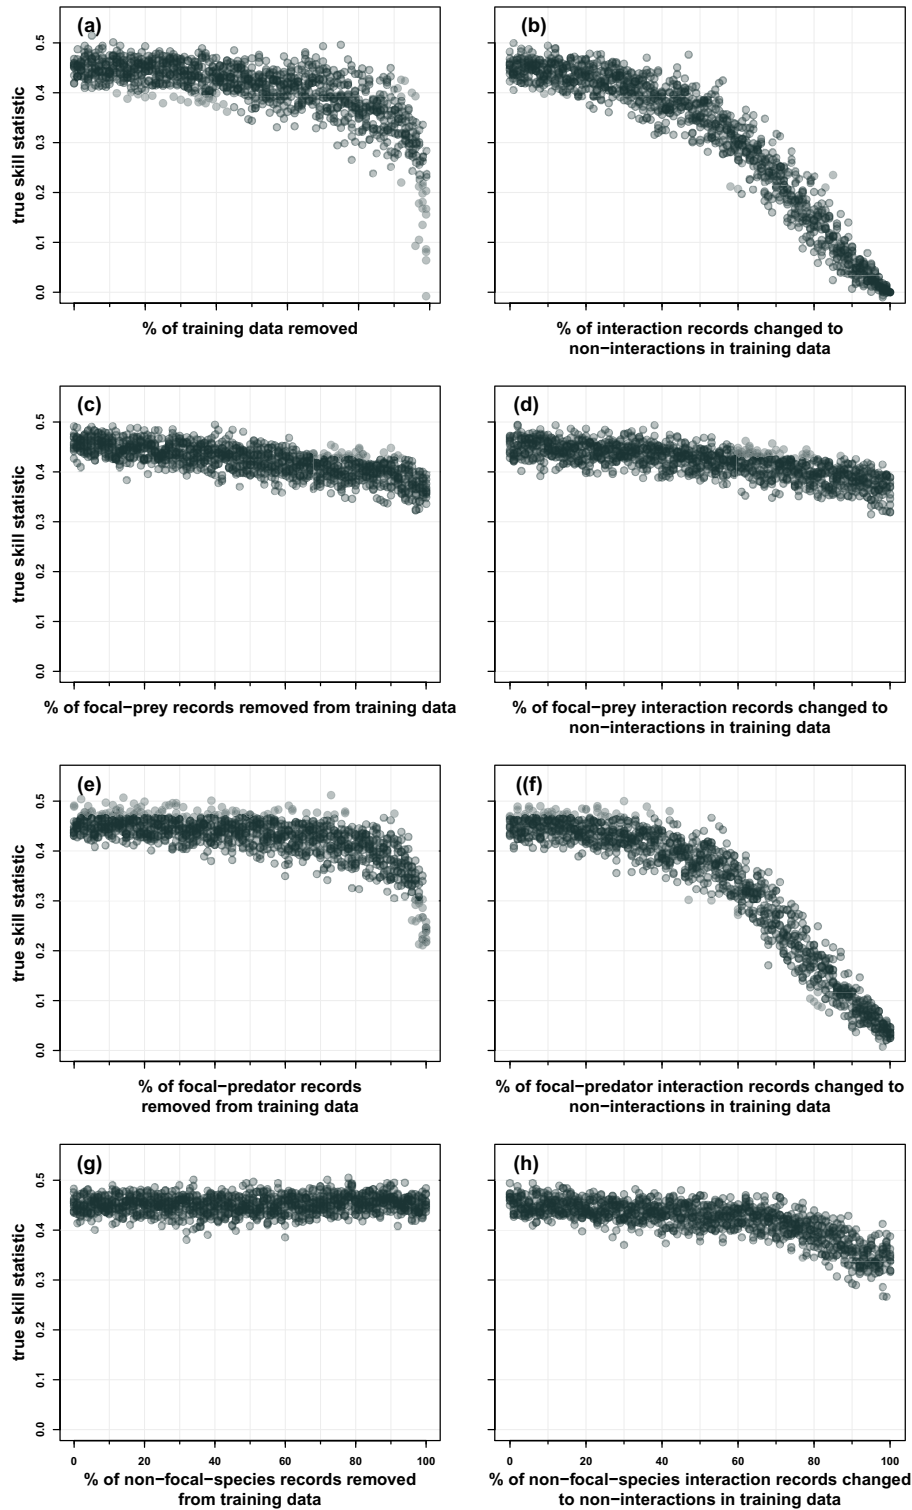


Figure 3. Effect of training-dataset quality on the performance of random forest for predicting predator–prey interactions in birds and mammals. Each panel shows performance when modifying a percentage of either the whole or specific taxonomic groups in the training dataset. (a) and (b) show performance when changing the whole dataset, (c) and (d) show performance when restricting changes to focal prey, (e) and (f) show performance when restricting changes to focal predators, and (g) and (h) show performance when restricting changes to non-focal (non-Simpson Desert) species. Panels in the left column show the effects of removing records from the training data, whereas panels in the right column show the effects of switching interaction records to non-interactions (false non-interactions). We trained random forests on the modified global predator–prey interaction dataset (i.e. restricted to predators with ≥ 5 prey records) and applied them to seven predators and their potential prey from the Simpson Desert. Plot points indicate performance (according to the true skill statistic) when applied to the Simpson Desert predators.




Figure 4. Effects of the quality of the training dataset on 1) the predicted relative suitability of prey and 2) the mean probabilities assigned to potential prey by random forests. Panels in the left column show the effects of reducing the number of interactions records involving the focal predator in the training dataset, whereas panels in the right column show the effect of switching interaction records involving the focal predator to non-interactions (false non-interactions). (a) and (b) show the correlation between prey probabilities assigned by the models trained on the changed datasets with that assigned by models trained on the complete (modified) training dataset. (c) and (d) show the mean probability (suitability) assigned to potential prey by models trained on the changed datasets minus those values assigned by the models trained on the complete (modified) training datasets. We trained these random forests on the modified global predator–prey interaction dataset (restricted to predators with ≥ 5 prey records) and applied them to seven predators and their potential prey from the Simpson Desert.

the probabilities assigned to potential prey were on average 0.17 ± 0.01 lower compared to those assigned by models trained on the full dataset.

Both the correlation and change in mean probability assigned to potential prey were more sensitive to false non-interactions involving the focal predators than to removing their records from the training dataset, showing more rapid declines (Fig. 4a versus b, c versus d). When changing all interactions involving the focal predators in the training data to non-interactions, the correlation between probabilities assigned to prey decreased to 0.27 ± 0.03 (Fig. 4b) and the mean probabilities assigned to potential prey dropped by 0.31 ± 0.01 (Fig. 4d).

Discussion

Our results demonstrate that random forests can predict the prey of predatory birds and mammals even when few or no interaction records involving the focal predator are included in training data (Fig. 4a). Applying random forests to predict predator–prey interactions for these predators (and

terrestrial vertebrates more generally) to facilitate network modelling of terrestrial vertebrate systems is therefore supported (Brousseau et al. 2018). However, our analyses highlight that caution is needed in applying random forest and other machine-learning approaches depending on the quality, quantity, and type of training data available, and in the face of these limitations we suggest that filtering training data could improve model performance (Fig. 5).

We found that random forests could predict bird and mammal interactions in both our global and ecosystem-specific datasets (including for native and recently introduced predators). Models based on ecomorphological, phylogenetic, or both types of traits, and those using few or many traits, could perform similarly well (Fig. 1) – consistent with previous studies showing that either data type can be used for predicting interactions in other taxa (Faisal et al. 2010, Morales-Castilla et al. 2015, Brousseau et al. 2018, Pomeranz et al. 2019, Elmasri et al. 2020) and that few traits/variables are required to make accurate predictions (Desjardins-Proulx et al. 2017). Because phylogenetic information is available for most terrestrial vertebrates (e.g. VertLife.org), ecological/morphological trait databases are

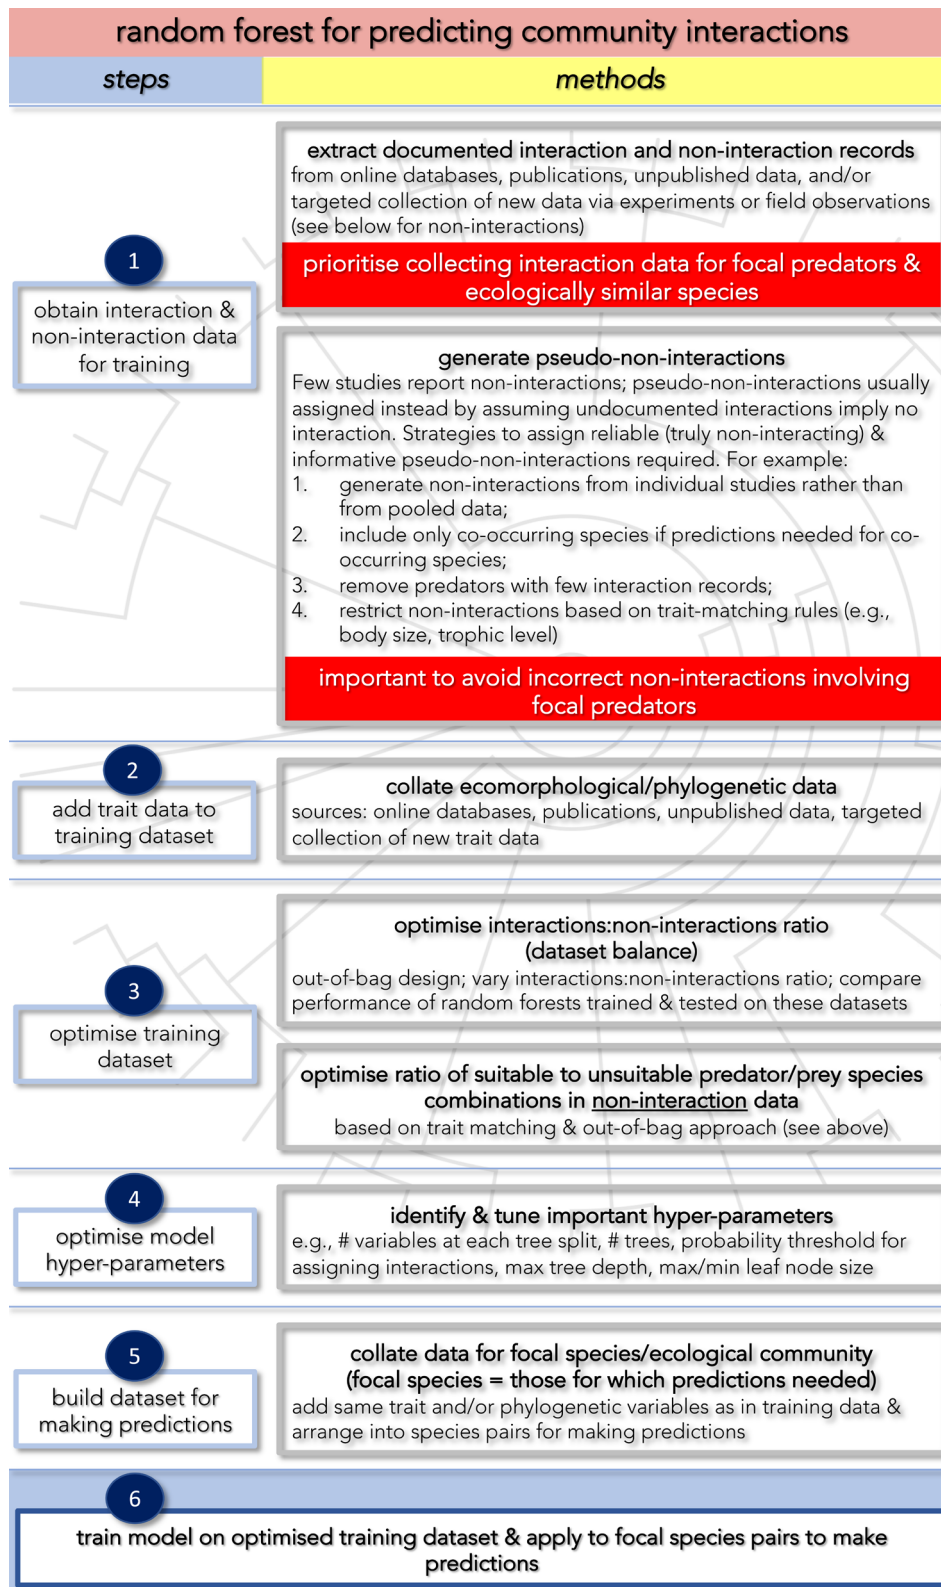


Figure 5. Framework for preparing and optimising random forest models for predicting species interactions. Most of these steps and methods are also relevant to the application of other machine learning approaches for inferring species interactions.

increasingly common and comprehensive (Jones et al. 2009, Oliveira et al. 2017, Herberstein et al. 2022), and interaction data for many species are available from open-source databases (e.g. globalbioticinteractions.org; web-of-life.es) and publications, random forests and other machine learning methods can now easily be applied to such data to predict predator–prey interactions for most vertebrate communities, and aid in the ecological network and invasive species modelling required to guide conservation decisions.

We found that model performance was robust to the random removal of records from the training dataset. In fact, we had to remove > 80% of records (> 15750 individual records) before an obvious change in performance occurred (Fig. 3a). Thus, models that use large training datasets are unlikely to be sensitive to random additions or removal of interaction and non-interaction records.

The removal of records involving focal predators only caused a prominent decline in model performance when we removed > 80% (Fig. 3e) of records in that group, but this translates to fewer (1240) records in total needing to be removed (compared to the random removal of records). Conversely, the complete removal of records involving focal-prey species or non-focal species had little impact on model performance (Fig. 3c, g). Thus, declines in model performance in the random-removal analyses were largely driven by the loss of focal predator records. This reduction in model performance following removal of focal-predator records was due to every prey species being assigned consistently lower probabilities and a slight shuffling of prey preference (Fig. 4a, c). Because relative prey preferences were largely maintained even when all training records involving the focal predator were removed from training data (mean Pearson $r=0.59$; Fig. 4a), our results suggest prey species could be predicted for predators that do not have any interaction data, which is likely the case for many understudied, rare, or extinct species. The ability to predict trophic interactions for extinct species accurately has important applications for understanding past extinction events (Ripple and Van Valkenburgh 2010, Pires et al. 2015, Llewelyn et al. 2022), while the capacity to predict interactions for rare species could help identify cryptic factors that threaten their persistence (Smith and Phillips 2006, Doherty et al. 2016). However, additional work to assess the requirements of training data for predicting interactions for species without interaction records is still required. For example, it is unclear how phylogenetically and/or ecologically similar species in the training data need to be to the focal species for models to be able to predict trophic interactions accurately.

In contrast to removal, false non-interactions in the training data reduced model performance more rapidly, triggering lower performance when > 40% of the interaction records were changed to non-interactions. This sensitivity highlights an important issue: most studies generate non-interactions by assuming that undocumented interactions imply true non-interaction; however, this assumption could generate many false non-interactions, especially if interaction records are sparse (i.e. if many of the predators in the training data have only a small proportion of their interactions documented).

To avoid generating false non-interactions, one can restrict non-interactions in training datasets to species combinations where the predator is too small to eat the potential prey and/or the potential predator comes from a low trophic level (Strona et al. 2021). However, applying this approach could reduce the capacity of random forests to identify unsuitable prey from within a predator's preferred prey-size range or for high trophic-level species. Indeed, the optimal ratio of non-interactions sampled from within vs outside preferred size range in the model that performed best on our Simpson Desert data (i.e. the many-variable, both data-types model) was 5:2 (Supporting information), suggesting that non-interactions of suitably-sized potential prey should be included in training data.

Alternatively, filtering training data by removing predators that have few interaction records could help avoid false non-interactions because many of the interactions/prey of these predators are likely missing (i.e. they have sparse interaction records, although these species would also include some predators with narrow diet breadths). When applied to the Simpson Desert data, our models trained on the global dataset modified by removing predators with < 5 interaction records performed better than those trained on the original/full global dataset (Fig. 3). This result might reflect a reduction of false non-interactions in the modified training dataset. Thus, removing predators/consumers that are missing most of their interactions from training datasets could be an effective way to filter training data and improve model performance.

Model performance was especially sensitive to false non-interactions in the training data involving focal predators. This type of change in data quality disrupted relative prey preferences and lowered probabilities assigned to interactions more extensively than did reducing the number of records involving focal predators (Fig. 4a versus b, c versus d), reinforcing the importance of avoiding false non-interactions – especially for focal predators – when generating pseudo-non-interactions for training data; otherwise, such data can change which species are identified as most-likely prey. Although removing predators/consumers with few interaction records from training data and/or restricting the generation of non-interactions based on relative size and trophic level are two simple methods for avoiding false non-interactions as we discussed, there are other strategies that could be applied to the entire- or focal-predator-component of training datasets to minimise the number of false non-interactions. For example, false non-interactions could be minimised by generating non-interactions on a case-by-case basis (similar to the 'target-group absences' approach used in species distribution modelling; Barbet-Massin et al. 2012), an approach shown to improve the accuracy of predator–prey interactions predicted by deep-learning models (Fricke et al. 2022). We generated non-interactions for the global dataset using interaction data from different regions, studies, and sources, but this approach could generate false non-interactions, such as when pairs of species are allopatric in studies that contributed to the training data and so are designated as

non-interacting, but the component species interact when/where they are sympatric. This approach could also generate misleading non-interactions, such as pairs of currently allopatric species that might interact should future conditions force them to become sympatric. Instead, training data non-interactions could be generated based on individual studies or locations of recorded interactions describing what a predator does and does not eat in a particular area. Such a nuanced approach would identify more reliable and informative non-interactions than would be achieved by assuming that any undocumented pairs from an incomplete global database are truly not predator and prey.

Our models performed better with the global dataset than with the Simpson Desert data. This difference might indicate that the random forests learnt patterns in the global dataset that are not as effective when re-applied to specific communities. For example, the random forests might have learnt from the global dataset that predators with particular traits do not eat prey with particular traits, but predators or prey with those traits might not be present in the Simpson Desert. In other words, learning the relationship did not help the random forest predict interactions in the Simpson Desert. Alternatively, random forests might learn which predators had few or many prey records within the global dataset – another pattern that would be less effective at identifying interactions and non-interactions when applied to the Simpson Desert predators due to the similarity among focal predators in terms of number of interactions in the training data (60% of predators in the original/full global dataset had < 5 interaction records whereas all Simpson Desert focal predators had > 10 interaction records in the global dataset). This might explain why our models did not perform as well on the global data when we used the modified dataset rather than the original/full global dataset for training and testing, because we had removed predators with few interaction records from the modified dataset. Geographic biases in the global training dataset could also lead to weaker performance on our ecosystem-specific dataset (Tuia et al. 2022). Irrespective of what caused the difference in model performance, our results indicate that performance based on global datasets might not be a strong indicator of performance when applied to single communities.

Previous studies applying random forests in aquatic and invertebrate communities have tended to report higher performance than our models achieved (mean true skill statistic consistently > 0.6; Desjardins-Proulx et al. 2017, Laigle et al. 2018, Brousseau et al. 2018, Parravicini et al. 2020, Strona et al. 2021). There are several possible explanations for this discrepancy, including previous studies 1) pooling prey species into broad resource categories (Parravicini et al. 2020), 2) using higher-quality training data (Laigle et al. 2018, Parravicini et al. 2020), 3) including resource predictions for a wider range of trophic groups, including easily-excluded non-interactions such as herbivores and/or detritivores eating other animals (Desjardins-Proulx et al. 2017), and 4) focusing on taxa whose interactions might be easier to predict than for terrestrial vertebrates (although further testing

is required to evaluate this possibility). Another potential means of improving the performance of random forests that neither we nor others have yet addressed is to include detail on intraspecific variation in traits and interactions, including ontogenetic shifts (González-Varo and Traveset 2016).

Random forest and other machine learning approaches are useful and flexible tools that have great potential in terms of predicting species interactions. While global interaction databases are a valuable resource for training and testing such models, they currently include only a small fraction of real-world species interactions. Furthermore, for the consumers that are included in such databases, in most cases only a small subset of their trophic interactions are documented, thereby increasing the likelihood of generating false non-interactions. Databases therefore need to be expanded, and a framework for generating non-interactions that avoids or minimises false non-interactions and retains informative non-interactions is required (Fig. 5 shows a proposed framework). By making such changes, the application of random forests and other machine learning methods (including ensemble approaches) can be further honed to infer species interactions more accurately.

Acknowledgements – We thank Phuong Chau and Tori Llewelyn for their helpful input. This work was supported by the Australian Research Council Centre of Excellence for Australian Biodiversity and Heritage (grant number CE170100015). We also acknowledge the Kurna, Gadigal and Wangkamadla people, the traditional custodians of the land on which this work was done.

Funding – This work was supported by the Australian Research Council Centre of Excellence for Australian Biodiversity and Heritage (grant no. CE170100015).

Author contributions

John Llewelyn: Conceptualization (equal); Data curation (lead); Formal analysis (lead); Methodology (lead); Project administration (lead); Validation (lead); Visualization (lead); Writing – original draft (lead); Writing – review and editing (equal). **Giovanni Strona:** Conceptualization (equal); Methodology (supporting); Writing – review and editing (equal). **Christopher Dickman:** Data curation (supporting); Writing – review and editing (equal). **Aaron Greenville:** Data curation (supporting); Writing – review and editing (equal). **Glenda Wardle:** Data curation (supporting); Writing – review and editing (equal). **Michael S. Y. Lee:** Data curation (supporting); Methodology (supporting); Writing – review and editing (equal). **Seamus Doherty:** Validation (supporting); Writing – review and editing (equal). **Farzin Shabani:** Methodology (supporting); Writing – review and editing (equal). **Frédéric Saltré:** Funding acquisition (supporting); Writing – review and editing (equal). **Corey Bradshaw:** Formal analysis (supporting); Funding acquisition (lead); Visualization (supporting); Writing – review and editing (equal).

Transparent peer review

The peer review history for this article is available at <https://publons.com/publon/10.1111/ecog.06619>.

Data availability statement

Data are available from the GitHub and Zenodo: <https://doi.org/10.5281/zenodo.7037189> (Llewelyn et al. 2023).

Supporting information

The Supporting information associated with this article is available with the online version.

References

- Allan, J. D., Castillo, M. M. and Capps, K. A. 2021. Energy flow and nutrient cycling in aquatic communities. – In: Castillo, M. M. and Capps, K. A., (eds), *Stream ecology: structure and function of running waters*. J. D. Springer International Publishing, pp. 357–381.
- Arif, S. and MacNeil, M. A. 2022. Predictive models aren't for causal inference. – *Ecol. Lett.* 25: 1741–1745.
- Barbet-Massin, M., Jiguet, F., Albert, C. H. and Thuiller, W. 2012. Selecting pseudo-absences for species distribution models: how, where and how many? – *Methods Ecol. Evol.* 3: 327–338.
- Blois, J. L., Zarnetske, P. L., Fitzpatrick, M. C. and Finnegan, S. 2013. Climate change and the past, present, and future of biotic interactions. – *Science* 341: 499–504.
- Bohan, D. A., Vacher, C., Tamaddoni-Nezhad, A., Raybould, A., Dumbrell, A. J. and Woodward, G. 2017. Next-generation global biomonitoring: large-scale, automated reconstruction of ecological networks. – *Trends Ecol. Evol.* 32: 477–487.
- Breiman, L. 2001. Random forests. – *Mach. Learn.* 45: 5–32.
- Brousseau, P.-M., Gravel, D. and Handa, I. T. 2018. Trait matching and phylogeny as predictors of predator–prey interactions involving ground beetles. – *Funct. Ecol.* 32: 192–202.
- Caron, D., Maiorano, L., Thuiller, W. and Pollock, L. J. 2022. Addressing the Eltonian shortfall with trait-based interaction models. – *Ecol. Lett.* 25: 889–899.
- Chamberlain, S., Barve, V., Mcglinn, D., Oldoni, D., Desmet, P., Geffert, L. and Ram, K. 2022. rgbif: interface to the global biodiversity information facility. – <https://CRAN.R-project.org/package=rgbif>
- Cox, N. et al. 2022. A global reptile assessment highlights shared conservation needs of tetrapods. – *Nature*: 285–290.
- Dellinger, A. S., Chartier, M., Fernández-Fernández, D., Penneys, D. S., Alvear, M., Almeda, F., Michelangeli, F. A., Staedler, Y., Armbruster, W. S. and Schönenberger, J. 2019. Beyond buzz-pollination – departures from an adaptive plateau lead to new pollination syndromes. – *New Phytol.* 221: 1136–1149.
- Desjardins-Proulx, P., Laigle, I., Poisot, T. and Gravel, D. 2017. Ecological interactions and the Netflix problem. – *PeerJ* 5: e3644.
- Doherty, T. S., Glen, A. S., Nimmo, D. G., Ritchie, E. G. and Dickman, C. R. 2016. Invasive predators and global biodiversity loss. – *Proc. Natl Acad. Sci. USA* 113: 11261–11265.
- Elmasri, M., Farrell, M. J., Davies, T. J. and Stephens, D. A. 2020. A hierarchical Bayesian model for predicting ecological interactions using scaled evolutionary relationships. – *Ann. Appl. Stat.* 14: 221–240.
- Faisal, A., Dondelinger, F., Husmeier, D. and Beale, C. M. 2010. Inferring species interaction networks from species abundance data: a comparative evaluation of various statistical and machine learning methods. – *Ecol. Inform.* 5: 451–464.
- Fleiss, J. L., Levin, B. and Paik, M. C. 2003. *Statistical methods for rates and proportions*. – John Wiley & Sons.
- Flinders University 2021. DeepThought (HPC). – Retrieved from <https://doi.org/10.25957/FLINDERS.HPC.DEEPThought>
- Fricke, E. C., Hsieh, C., Middleton, O., Gorkzynski, D., Cappello, C. D., Sanisidro, O., Rowan, J., Svenning, J. C. and Beaudrot, L. 2022. Collapse of terrestrial mammal food webs since the Late Pleistocene. – *Science* 377: 1008–1011.
- González-Varo, J. P. and Traveset, A. 2016. The labile limits of forbidden interactions. – *Trends Ecol. Evol.* 31: 700–710.
- Gravel, D., Poisot, T., Albouy, C., Velez, L. and Mouillot, D. 2013. Inferring food web structure from predator–prey body size relationships. – *Methods Ecol. Evol.* 4: 1083–1090.
- Guénard, G., Legendre, P. and Peres-Neto, P. 2013. Phylogenetic eigenvector maps: a framework to model and predict species traits. – *Methods Ecol. Evol.* 4: 1120–1131.
- Guénard, G. and Legendre, P. 2022. MPSEM: modeling phylogenetic signals using eigenvector maps. – <https://CRAN.R-project.org/package=MPSEM>
- Herberstein, M. E., McLean, D. J., Lowe, E., Wolff, J. O., Khan, M. K., Smith, K., Allen, A. P., Bulbert, M., Buzatto, B. A., Eldridge, M. D. B., Falster, D., Fernandez Winzer, L., Griffith, S. C., Madin, J. S., Narendra, A., Westoby, M., Whiting, M. J., Wright, I. J. and Carthey, A. J. R. 2022. AnimalTraits – a curated animal trait database for body mass, metabolic rate and brain size. – *Sci. Data* 9: 265.
- Hortal, J., de Bello, F., Diniz-Filho, J. A. F., Lewinsohn, T. M., Lobo, J. M. and Ladle, R. J. 2015. Seven shortfalls that beset large-scale knowledge of biodiversity. – *Annu. Rev. Ecol. Evol. Syst.* 46: 523–549.
- Jones, K. E. et al. 2009. PanTHERIA: a species-level database of life history, ecology, and geography of extant and recently extinct mammals. – *Ecol. Arch.* 90: 2648–2648.
- Jordano, P. 2016. Chasing ecological interactions. – *PLoS Biol.* 14: e1002559.
- Kaneryd, L., Borrvall, C., Berg, S., Curtsdotter, A., Eklöf, A., Hauzy, C., Jonsson, T., Münger, P., Setzer, M., Säterberg, T. and Ebenman, B. 2012. Species-rich ecosystems are vulnerable to cascading extinctions in an increasingly variable world: extinction cascades in a variable world. – *Ecol. Evol.* 2: 858–874.
- Klomborg, Y., Tropek, R., Mertens, J. E. J., Kobe, I. N., Hodeček, J., Raška, J., Fominka, N. T., Souto-Vilarós, D., Janečková, P. and Janeček, Š. 2022. Spatiotemporal variation in the role of floral traits in shaping tropical plant–pollinator interactions. – *Ecol. Lett.* 25: 839–850.
- Kotula, H. J., Peralta, G., Frost, C. M., Todd, J. H. and Tylianakis, J. M. 2021. Predicting direct and indirect non-target impacts of biocontrol agents using machine-learning approaches. – *PLoS One* 16: e0252448.
- Laigle, I., Aubin, I., Digel, C., Brose, U., Boulangéat, I. and Gravel, D. 2018. Species traits as drivers of food web structure. – *Oikos* 127: 316–326.
- Liem, K. F. 1984. Functional versatility, speciation, and niche overlap: are fishes different? – Routledge.
- Liem, K. F. 1990. Aquatic versus terrestrial feeding modes: possible impacts on the trophic ecology of vertebrates 1. – *Am. Zool.* 30: 209–221.
- Llewelyn, J. 2022. Random forests for predicting predator–prey interactions in terrestrial vertebrates. – GitHub and Zenodo 1. <https://github.com/JohnLlewelyn/random-forests-for-predicting-predator-prey-interactions-in-terrestrial-vertebrates/releases/tag/V1>.

- Llewelyn, J., Strona, G., McDowell, M. C., Johnson, C. N., Peters, K. J., Stouffer, D. B., de Visser, S. N., Saltré, F. and Bradshaw, C. J. A. 2022. Sahul's megafauna were vulnerable to plant-community changes due to their position in the trophic network. – *Ecography* 2022: e06089.
- Llewelyn, J., Strona, G., Dickman, C. R., Greenville, A. C., Wardle, G. M., Lee, M. S. Y., Doherty, S., Shabani, F., Saltré, F. and Bradshaw, C. J. A. 2023. Data from: Predicting predator–prey interactions in terrestrial endotherms using random forest. – Zenodo Digital Repository, <https://doi.org/10.5281/zenodo.7037189>.
- Malley, J. D., Kruppa, J., Dasgupta, A., Malley, K. G. and Ziegler, A. 2012. Probability machines: consistent probability estimation using nonparametric learning machines. – *Methods Inf. Med.* 51: 74–81.
- Mateo, R. G., Croat, T. B., Felicísimo, Á. M. and Muñoz, J. 2010. Profile or group discriminative techniques? Generating reliable species distribution models using pseudo-absences and target-group absences from natural history collections: reliability of models using pseudo-absences and target-group absences. – *Divers. Distrib.* 16: 84–94.
- Miller-ter Kuile, A. et al. 2022. Predator–prey interactions of terrestrial invertebrates are determined by predator body size and species identity. – *Ecology* 103: e3634.
- Morales-Castilla, I., Matias, M. G., Gravel, D. and Araújo, M. B. 2015. Inferring biotic interactions from proxies. – *Trends Ecol. Evol.* 30: 347–356.
- Morton, D. N., Keyes, A., Barner, A. K. and Dee, L. E. 2022. Merging theory and experiments to predict and understand coextinctions. – *Trends Ecol. Evol.* 37: 886–898.
- Oliveira, B. F., São-Pedro, V. A., Santos-Barrera, G., Penone, C. and Costa, G. C. 2017. AmphiBIO, a global database for amphibian ecological traits. – *Sci. Data* 4: 170123.
- Ornaí, A. and Keasar, T. 2020. Floral complexity traits as predictors of plant–bee interactions in a Mediterranean pollination web. – *Plants (Basel)* 9: 1432.
- Paradis, E. and Schliep, K. 2019. ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. – *Bioinformatics* 35: 526–528.
- Parravicini, V., Casey, J. M., Schiettekatte, N. M. D., Brandl, S. J., Pozas-Schacre, C., Carlot, J., Edgar, G. J., Graham, N. A. J., Harmelin-Vivien, M., Kulbicki, M., Strona, G. and Stuart-Smith, R. D. 2020. Delineating reef fish trophic guilds with global gut content data synthesis and phylogeny. – *PLoS Biol.* 18: e3000702.
- Pichler, M., Boreux, V., Klein, A.-M., Schleuning, M. and Hartig, F. 2020. Machine learning algorithms to infer trait-matching and predict species interactions in ecological networks. – *Methods Ecol. Evol.* 11: 281–293.
- Pires, M. M., Koch, P. L., Fariña, R. A., de Aguiar, M. A. M., dos Reis, S. F. and Guimarães, P. R. 2015. Pleistocene megafaunal interaction networks became more vulnerable after human arrival. – *Proc. R. Soc. B* 282: 20151367.
- Poelen, J. H., Simons, J. D. and Mungall, C. J. 2014. Global Biotic Interactions: an open infrastructure to share and analyze species-interaction datasets. – *Ecol. Inform.* 24: 148–159.
- Poisot, T. 2023. Guidelines for the prediction of species interactions through binary classification. – *Methods Ecol. Evol.* 14: 1333–1345.
- Pomeranz, J. P. F., Thompson, R. M., Poisot, T. and Harding, J. S. 2019. Inferring predator–prey interactions in food webs. – *Methods Ecol. Evol.* 10: 356–367.
- Pringle, R. M. and Hutchinson, M. C. 2020. Resolving food-web structure. – *Annu. Rev. Ecol. Evol. Syst.* 51: <https://doi.org/10.1146/annurev-ecolsys-110218-024908>.
- Ripple, W. J. and Van Valkenburgh, B. 2010. Linking top-down forces to the Pleistocene megafaunal extinctions. – *J. Biol. Sci.* 60: 516–526.
- Samson, F. B. and Knopf, F. L. 1996. Ecosystem management: selected readings. – Springer.
- Shurin, J. B., Gruner, D. S. and Hillebrand, H. 2006. All wet or dried up? Real differences between aquatic and terrestrial food webs. – *Proc. R. Soc. B* 273: 20053377.
- Sinclair, A. R. E., Mduma, S. and Brashares, J. S. 2003. Patterns of predation in a diverse predator–prey system. – *Nature* 425: 288–290.
- Smith, J. G. and Phillips, B. L. 2006. Toxic tucker: the potential impact of cane toads on Australian reptiles. – *Pac. Conserv. Biol.* 12: 40–49.
- Strona, G. 2022. Hidden pathways to extinction. – Springer International Publishing.
- Strona, G. and Bradshaw, C. J. A. 2018. Co-extinctions annihilate planetary life during extreme environmental change. – *Sci. Rep.* 8: 16724.
- Strona, G. and Bradshaw, C. J. A. 2022. Coextinctions dominate future vertebrate losses from climate and land use change. – *Sci. Adv.* 8: eabn4345.
- Strona, G., Beck, P. S. A., Cabeza, M., Fattorini, S., Guilhaumon, F., Micheli, F., Montano, S., Ovaskainen, O., Planes, S., Veech, J. A. and Parravicini, V. 2021. Ecological dependencies make remote reef fish communities most vulnerable to coral loss. – *Nat. Commun.* 12: 7282.
- Strydom, T., Catchen, M. D., Banville, F., Caron, D., Dansereau, G., Desjardins-Proulx, P., Forero-Muñoz, N. R., Higinio, G., Mercier, B., Gonzalez, A., Gravel, D., Pollock, L. and Poisot, T. 2021. A roadmap towards predicting species interaction networks (across space and time). – *Phil. Trans. R. Soc. B* 376: 20210063.
- Strydom, T., Bouskila, S., Banville, F., Barros, C., Caron, D., Farrell, M. J., Fortin, M.-J., Hemming, V., Mercier, B., Pollock, L. J., Runghen, R., Dalla Riva, G. V. and Poisot, T. 2022. Food web reconstruction through phylogenetic transfer of low-rank network representation. – *Methods Ecol. Evol.* 13: 2838–2849.
- Sydenham, M. A. K., Venter, Z. S., Reitan, T., Rasmussen, C., Skrindo, A. B., Skoog, D. I. J., Hanevik, K.-A., Hegland, S. J., Dupont, Y. L., Nielsen, A., Chipperfield, J. and Rusch, G. M. 2021. MetaComNet: a random forest-based framework for making spatial predictions of plant–pollinator interactions. – *Methods Ecol. Evol.* 13: 500–513.
- Thompson, J. N. 1999. The evolution of species interactions. – *Science* 284: 2116–2118.
- Tilman, D., Clark, M., Williams, D. R., Kimmel, K., Polasky, S. and Packer, C. 2017. Future threats to biodiversity and pathways to their prevention. – *Nature* 546: 73–81.
- Traill, L. W., Lim, M. L. M., Sodhi, N. S. and Bradshaw, C. J. A. 2010. Mechanisms driving change: altered species interactions and ecosystem function through global warming. – *J. Anim. Ecol.* 79: 937–947.
- Tuía, D., Kellenberger, B., Beery, S., Costelloe, B. R., Zuffi, S., Risse, B., Mathis, A., Mathis, M. W., van Langevelde, F., Burghardt, T., Kays, R., Klinck, H., Wikelski, M., Couzin, I. D., van Horn, G., Crofoot, M. C., Stewart, C. V. and Berger-Wolf, T. 2022. Perspectives in machine learning for wildlife conservation. – *Nat. Commun.* 13: 792.

- Tylianakis, J. M., Didham, R. K., Bascompte, J. and Wardle, D. A. 2008. Global change and species interactions in terrestrial ecosystems. – *Ecol. Lett.* 11: 1351–1363.
- Valdovinos, F. S., Berlow, E. L., Moisset de Espanés, P., Ramos-Jiliberto, R., Vázquez, D. P. and Martínez, N. D. 2018. Species traits and network structure predict the success and impacts of pollinator invasions. – *Nat. Commun.* 9: 2153.
- Van der Putten, W. H., Macel, M. and Visser, M. E. 2010. Predicting species distribution and abundance responses to climate change: why it is essential to include biotic interactions across trophic levels. – *Phil. Trans. R. Soc. B Biol. Sci.* 365: 2025–2034.
- Wilman, H., Belmaker, J., Simpson, J., de la Rosa, C., Rivadeneira, M. M. and Jetz, W. 2014. EltonTraits 1.0: species-level foraging attributes of the world's birds and mammals. – *Ecology* 95: 2027–2027.
- Wright, M. N. and Ziegler, A. 2017. ranger: a fast implementation of Random Forests for high dimensional data in C++ and R. – *J. Stat. Soft.* 77: 1–17.