Quantitative Prediction of Interactions in Bipartite Networks Based on Traits, Abundances, and Phylogeny*

Gita Benadi,[1,†] Carsten F. Dormann,[1,2,†,‡] Jochen Fründ,[1] Ruth Stephan,[1] and Diego P. Vázquez[2,3]

1. Biometry and Environmental System Analysis, University of Freiburg, Freiburg, Germany;   2. Freiburg Institute of Advanced Studies (FRIAS), University of Freiburg, Freiburg, Germany;   3. Argentine Institute for Dryland Research, Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Mendoza, Argentina; and Faculty of Exact and Natural Sciences, National University of Cuyo, Mendoza, Argentina

ABSTRACT: Ecological interactions link species in networks. Loss of species from or introduction of new species into an existing network may have substantial effects for interaction patterns. Predicting changes in interaction frequency while allowing for rewiring of existing interactions—and hence estimating the consequences of community compositional changes—is thus a central challenge for network ecology. Interactions between species groups, such as pollinators and flowers or parasitoids and hosts, are moderated by matching morphological traits or sensory clues, most of which are unknown to us. If these traits are phylogenetically conserved, however, we can use phylogenetic distances to construct latent, surrogate traits and try to match those across groups, in addition to observed traits. Understanding how important traits and trait matching are, relative to abundances and chance, is crucial to estimating the fundamental predictability of network interactions. Here, we present a statistically sound approach ("tapnet") to fitting abundances, traits, and phylogeny to observed network data to predict interaction frequencies. We thereby expand existing approaches to quantitative bipartite networks, which so far have failed to correctly represent the nonindependence of network interactions. Furthermore, we use simulations and cross validation on independent data to evaluate the predictive power of the fit. Our results show that tapnet is on a par with abundance-only, matching centrality, and machine learning approaches. This approach also allows us to evaluate how well current concepts of trait matching work. On the basis of our results, we expect that interactions in well-sampled networks can be well predicted if traits and abundances are the main driver of interaction frequency.

Introduction

Network ecology has vastly increased our knowledge of ecosystems, delivered fascinating insights into their organization (e.g., the slow and fast energy channels in Rooney et al. 2006), and fostered speculation about coevolution (e.g., Guimarães et al. 2011). While arguably the litmus test of any ecological understanding is a test of its predictive power (Houlahan et al. 2017), only recently have predictive models of network interactions started to emerge (e.g., Ives and Godfray 2006; Petchey et al. 2008; Crea et al. 2016; Rohr et al. 2016; Brousseau et al. 2018; Pichler et al. 2020). After several decades of research on ecological networks we have learned a lot about their patterns (e.g., along latitudinal or elevational gradients, across different types of interactions) and their stability (see, e.g., Bascompte and Jordano 2014; Moore et al. 2017), yet so far we have had limited success in predicting interaction intensities or even aggregate network structure (as pointed out by Vázquez et al. 2009; Olito and Fox 2015; Poisot et al. 2016; Valdovinos et al. 2018; but see Pomeranz et al. 2019; Vizentin-Bugoni et al. 2020).

Some of the most successful attempts at predicting network structure were made in studies of food webs, where body size ratios were found to be a good predictor of the presence or absence of predator-prey interactions (e.g., Allesina 2011; Gravel et al. 2013; Pomeranz et al. 2019) and allometric scaling of parameters allowed a mechanistic model of optimal foraging theory (the contingency model;

Stephens and Krebs 1986) to be fitted to large networks (Beckerman et al. 2006; Petchey et al. 2008). However, body size is not a suitable predictor for most other interaction types, such as host-parasitoid interactions or interactions involving plants (Bascompte and Jordano 2014; Dormann and Blüthgen 2017). Moreover, while network predictions should ideally be based on theory, existing theoretical models such as the contingency model make strongly simplifying assumptions (e.g., sequential encounter of resources, maximizing of average resource intake, ideal knowledge of resource availability, no within-guild interference; Pyke 1984; Stephens and Krebs 1986) but still require large numbers of parameter values as input, which are difficult to provide without a shortcut, such as allometric scaling. Thus, we still lack a general approach to predicting various types of ecological networks. In addition, whereas most existing models of ecological networks only predict binary network structure, models that allow estimation of the intensity of interspecific interactions can glean more information from the data.

While a theoretical model encompassing all types of interaction networks does not yet exist, there are obvious candidate factors for predicting interaction intensities, some at the level of the individual species (such as their abundance), some at the interplay of species (such as the match of interaction-relevant traits; Junker et al. 2010, 2013; Dehling et al. 2016). Since some of the relevant traits are difficult to measure but show a phylogenetic signal, phylogeny can be used as a proxy for such unmeasured traits (Ives and Godfray 2006; Pearse and Altermatt 2013; Morales-Castilla et al. 2015; Peralta 2016).

There have been several previous approaches to including traits, abundance, and phylogeny in a statistical method to analyze and possibly predict interactions. They differ in their scope (binary vs. quantitative networks), the type of network (bipartite vs. unipartite), and the statistical method. For example, Ives and Godfray (2006), Pearse and Altermatt (2013), Rohr et al. (2016), and Crea et al. (2016) all used regression models to predict binary networks based on phylogeny and measured traits. Brousseau et al. (2018) improved on the model of Rohr et al. (2016) by adding a larger number of traits as predictors and using a more flexible generalized additive model (GAM). Vázquez et al. (2009) and Olito and Fox (2015) chose a different statistical approach in which they constructed matrices of interaction probabilities based on traits and abundances and assumed that the observed interaction intensities are drawn from a multinomial distribution with these probabilities. Most recently, Desjardins-Proulx et al. (2017) and Pichler et al. (2020) modeled interaction networks using machine learning algorithms, which are highly flexible but provide little information on the underlying mechanisms of the interactions.

While the above-mentioned approaches have had some success in describing ecological networks, they face two issues, one statistical, the other ecological. First, most methods treat the interactions in a network as statistically independent (except Vázquez et al. 2009; Olito and Fox 2015; Crea et al. 2016), although this assumption is likely to be violated. Any interaction with one species precludes interaction with other species at the same time; thus, one more observation here inevitably means one less there. Furthermore, depending on the type of interaction and method of data collection, the same individual may be observed multiple times, again violating the independence assumption. Finally, interactions of different consumer species may be nonindependent because of intraguild competition for resources, which can cause shifts in species' preferences (e.g., Loeuille and Loreau 2005; Spiesman and Gratton 2016). Such nonindependence must be accounted for in order to not yield biased and overconfident model estimates.

An ecological issue is how to represent the role of species traits and their matching. Models based on linear regression assume that all traits and trait combinations have linear effects on interaction probabilities, while machine learning algorithms do not provide any information on the mechanisms connecting traits to interactions. In reality, interaction intensities may depend on the matching of quantitative trait values in nonlinear and possibly asymmetric ways. For instance, large-billed birds can feed on small seeds, but small-billed ones cannot feed on large seeds (e.g., Muñoz et al. 2017). Ideally, models of ecological networks should allow explicit incorporation of such trait-matching mechanisms, to correctly represent the ecological mechanisms and provide accurate predictions (but for an implicit approach, see Sebastián-González et al. 2016).

In this article, we present a statistical approach to analyzing and predicting interaction intensity based on observed and phylogeny-based latent traits and their matching, alongside abundances of each species. Unlike previous analyses, which were mostly explorative, we assess our model's performance on independent network data that were not used for fitting. For more ecological realism and as an extension of previous approaches, we provide symmetric and asymmetric trait-matching functions. In contrast to all previous approaches for traits and phylogeny, we account for the nonindependence of observations in the network using the multinomial probability approach of Vázquez et al. (2009) and Olito and Fox (2015). Finally, we assess the quality of our approach at the level of the individual link, not only at the level of network patterns, as the aim is to predict a specific interaction.

Our approach can be used for prediction of, say, introduction or loss of species from a community, which may lead to a "rewiring" of interactions, or for predicting the

effects of changes in abundances. Moreover, it can be used to quantify the importance of observed trait pairs for such predictions, such as the morphological compatibility of a plant's flower corolla and a pollinator's proboscis or the beak size of a frugivorous bird and the size of a fruit (see also Pichler et al. 2020).

We first outline the idea and details behind our approach, then demonstrate its validity with simulated data, and finally use three hummingbird-flower networks to make predictions across habitats. In that case study, we also use alternative published approaches to gauge their performances on real-world validation data.

Methods

We call our approach "tapnet," as it uses traits, abundance activity, and phylogeny to predict network interactions. Our background in pollination ecology makes it natural for us to think of bipartite interaction networks, where one group's members (e.g., pollinators) interact with mem-

bers of another group (e.g., plants), but not within each group. However, the approach can be similarly applied to functional group- or individual-based networks and probably extended to one-mode networks, but that is beyond the scope of our study. In the case of analyzing individuals, species' average traits would be replaced by individual trait values. Abundances could either be removed completely or be replaced by some measure of individual activity.

Before explaining our approach in detail, here is an outline (fig. 1). We developed a model that outputs expected interaction probabilities based on traits, abundances, and phylogenies. We can now compare the output with an observed network and optimize model parameters so as to maximize fit. In this way, we estimate several (largely) ecologically interpretable parameters. For prediction, we can use the fitted model together with new abundances (including previously unobserved species) and yield expected interaction probabilities. While the role of traits is probably clear, phylogenies are used to construct so-called latent traits to be matched across groups. As a side effect, new
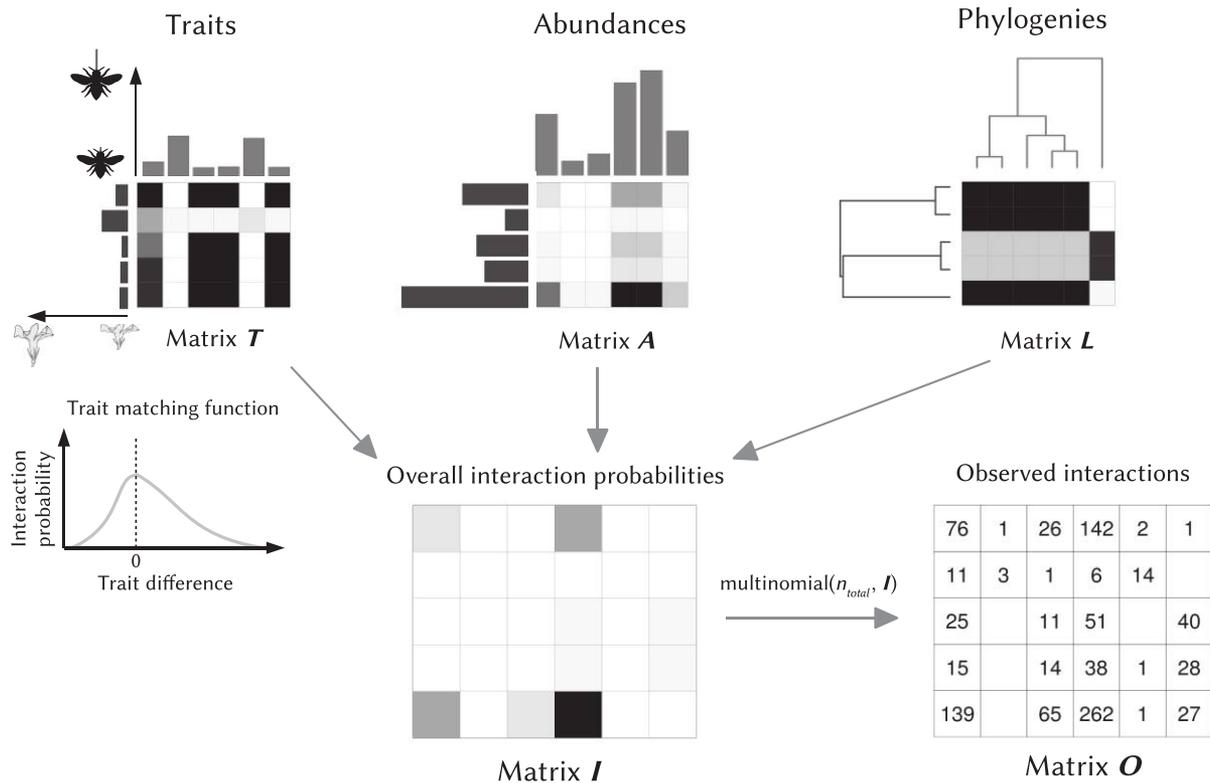

Figure 1: Conceptual overview of the tapnet approach. Tapnet combines information on trait matching (**T**), abundances (**A**), and phylogeny-derived matching latent traits (**L**; *top row*) into a single matrix **I** of predicted interaction probabilities. The observed interaction network **O** is assumed to be drawn from a multinomial distribution with probabilities given by **I** and total number of interactions n_{total} equaling the observed number of interactions. Traits have to be provided in pairs across the two levels, which are then compared using a possibly asymmetric trait-matching function. Each input yields an independent matrix of interaction probabilities, which are then multiplied and rescaled to yield the final interaction matrix **I**. Matrix cells are filled with different shades of gray to indicate interaction probabilities from 0 (white) to 1 (black). Some values that are close to zero appear white.

species entering a community can be positioned in the phylogeny, and this position then feeds through to the actual predictions, in addition to the (optional) observed traits.

Traits refer to species-specific characteristics that have a counterpart in a trait of the other group, and the analyst must provide these traits in matching format (i.e., corolla and proboscis length, or phenologies during a season, quantified in such a way that the same value for both would imply a perfect match). The model can be employed without trait information.

Abundance activity (henceforth, "abundances") can be any measure proportional to the probability of encountering a species in the field. For instance, a plant with an attractive scent would be more conspicuous to a pollinator than a nonodorant one, and an abundant bee would be encountered more often than a rare one if they move at the same rate (hence, abundance activity). If the trait pairs responsible for interaction activity are known and provided to tapnet, the role of actual abundances can be quantified; otherwise, the two will remain confounded (as in any other analysis).

Phylogenies for each group are used as building blocks for unobserved (latent) traits moderating species interactions. They add complementary information to the observed traits (Pearse et al. 2013) and may be able to suggest possible additional traits that underlie the observed interactions.

We present the approach starting with the statistical goal and then going into ever more detail. We provide R code for simulating tapnet data, for fitting observed networks, for assessing model fit, and for predicting new abundances and traits in the R package tapnet.[1]

Likelihood

Our approach is to fit a (nonstandard) statistical model to observed data on interaction networks. While we can use several networks simultaneously to fit the model, we restrict our outline here to the simplest case of only one observed network. Let us call the observed interaction matrix $\mathbf{O}$ of dimension $m \times n$ (m rows by n columns).

We compare $\mathbf{O}$ with our model predictions $\mathbf{P}$ by means of a multinomial distribution (eq. [1]), as entries in that interaction matrix are nonindependent. (Entries in $\mathbf{O}$ are integers, typically the number of observed interactions per standardized observation effort. For continuously valued observations, a Dirichlet distribution could probably be employed; Crea et al. 2016.) To do so, we rep-

resent $\mathbf{O}$ as a vector of length mn rather than a matrix. Formally, we predict the entries in $\mathbf{O}$, o_i, on the basis of the total number of observed interactions, $n_{\text{total}} = \sum_{i=1}^{mn} o_i$, and the vector of predicted interaction probabilities, $\hat{p}_i$, which are the output of our model:

$$f(o_1, \ldots, o_{mn}; n_{\text{total}}, \hat{p}_1, \ldots, \hat{p}_{mn}) = \frac{n_{\text{total}}!}{\prod_i o_i!} \prod_{i=1}^{mn} \hat{p}_i^{o_i}. \quad (1)$$

Model Components

Our model yields a matrix of predicted interactions, $\mathbf{P} = (\hat{p}_{i,j}) \in \mathbb{R}^{m \times n}$, as a (rescaled) Hadamard ($=$ element-wise) product of three prediction components—abundance-based expected probabilities $\mathbf{A}$, trait matching–based expected probabilities $\mathbf{T}$, and latent trait–based expected probabilities $\mathbf{L}$—each scaled to sum to 1:

$$\mathbf{P} = \mathbf{A} \circ \frac{(\mathbf{T} \circ \mathbf{L})^\delta}{\sum (\mathbf{T} \circ \mathbf{L})^\delta}. \quad (2)$$

The term $\mathbf{A}$ is the matrix of abundance-based interaction probabilities based on the cross product of normalized species abundances vectors for the lower o_A^l and higher o_A^h trophic level (note that we reserve capital letters for matrices): $\mathbf{A} = (o_A^l / \sum o_A^l) \times (o_A^h / \sum o_A^h)$. Thus, $\mathbf{A}$ represents the probability of an interaction for each cell based only on the relative abundances of the different species; we would expect more interactions among common species than among rare species. The term $o_A^{l,h}$ is based on independently measured abundances, not simply the marginal totals of the observed matrix $\mathbf{O}$. The free exponent, $\delta \in (0, 1)$, allows the optimization to give more or less weight to traits relative to abundance and thereby also serves as a quantification of the importance of abundance within the fit. As a side effect, the abundance-only predictions are a limiting case of tapnet, for $\delta = 0$. The denominator summing over $(\mathbf{T} \circ \mathbf{L})^\delta$ is necessary to renormalize this term to sum to 1, on par with $\mathbf{A}$.

The term $\mathbf{T}$ is the matrix of expected interaction probabilities based on the degree to which observed traits match between species of the different groups. If, say, the proboscis of a pollinator is a bit too short or too long compared with the depth of the corolla of a plant species, then interactions become less likely than perfectly matching lengths. We define a (single parameter) trait-matching function for a pair (i, j) of trait values $f_t(t_i^l, t_j^h)$ by the Gaussian function:

$$f_t(t_i^l, t_j^h, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{\left(t_j^h - t_i^l\right)^2}{2\sigma^2}\right). \quad (3)$$

Alternatively, we can define an asymmetric matching function (a lognormal with its mode shifted to zero), as

---

1. Code that appears in *The American Naturalist* is provided as a convenience to readers. It has not necessarily been tested as part of peer review.

too long a proboscis is no obstacle to an interaction, while one too short is

$$f_t(t_i^l, t_j^h, \mu) =$$

$$\frac{1}{(t_j^h - t_i^l + e^{\mu-1})\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\ln(t_j^h - t_i^l + e^{\mu-1}) - \mu)^2\right),$$

(4)

where $e^{\mu-1}$ is a parameter determining the shape of the function, in a way similar to σ in equation (3). It is fit during the model optimization (see below). Other functions could of course be used, especially when the relationship between trait values and interaction probability is known. The trait-matching function is computed for all species; that is, the inputs are in fact vectors of trait values for the lower ($\mathbf{t}^l$) and higher ($\mathbf{t}^h$) level. The elements of $\mathbf{T}$, (τ_{ij}), are the f_t values computed using equation (3) or (4): (τ_{ij}) = $f_t(l_i^l, l_j^h)$.

Latent trait–based interaction probability matrix $\mathbf{L}$, finally, is the most complex matrix of expected interaction probabilities. Its function is primarily to improve predictions, as abundances and observed traits are unlikely to be sufficient to capture the information present in the observed interaction matrix. Just like $\mathbf{T}$, $\mathbf{L}$ is also found by matching traits of one level to those of the other by optimizing the σ parameter of its Gaussian trait-matching function (eq. [3]). In this case, however, these traits are unobserved (latent) and are constructed as part of the optimization process from phylogenetic eigenvectors, as explained in the next section.

Constructing Matching Latent Trait Values

The reasons why a species from one group interacts with one from another may be extremely complex. Bees, for example, may rely on visual and olfactorial signals, as well as innate preferences and learned behavior, for selecting a flower to visit (Chittka and Raine 2006). On the other hand, plants may increase nectar production in response to sensing pollinators' wingbeats (Veits et al. 2019). If these traits can be measured, they can be used in tapnet as observed traits. In many cases, however, we may have little chance to identify, let alone measure, the traits involved. In such cases, one can "invent" trait values with the sole aim of improving the fit of the model.

Computationally, it is more efficient—and ecologically more satisfactory—if there is a basis for computing such latent trait values (similar to geographic distances in spatial models used in joint species distribution models; Warton et al. 2015). In our case, we assume latent traits to summarize traits correlated with phylogeny. These latent traits typically exhibit no phylogenetic signal anymore, and the phylogenies are used only as an efficient way to generate

orthogonal vectors. This approach has several benefits. First, it allows us to introduce new species when predicting from the fitted tapnet model, as we can compute their phylogenetic position relative to the other species in the group and hence also the value of their latent trait. Second, the resulting latent trait may, on mapping it to the phylogeny, suggest a hypothesis about an actual trait behind it.

Technically, constructing a latent trait within a group is straightforward if a phylogeny is available. Alternatively, a taxonomy can be used (Clarke and Warwick 1999). From such a phylogeny one can compute phylogenetic eigenvectors (Guénard et al. 2013), that is, for k species $k-1$ vectors that are orthogonal to each other and represent the information on the phylogeny (similar to the way a principal component analysis summarizes the information in a data set): the first phylogenetic eigenvector accounts for the largest genetic differences in the tree, the second accounts for the largest differences in the remaining variation, and so forth. Mathematically, a phylogenetic tree can be represented as a (cophenetic) distance matrix, and an eigenvalue decomposition of said distance matrix yields the phylogenetic eigenvectors.

For each group, we can now define a latent trait vector $\mathbf{l}$ as a linear combination of the phylogenetic eigenvectors $\mathbf{p}_i, \ldots, \mathbf{p}_{k-1}$, one for each trophic level:

$$\mathbf{l}^l = a_1\mathbf{p}_1^l + a_2\mathbf{p}_2^l + a_3\mathbf{p}_3^l + \cdots + a_{m-1}\mathbf{p}_{m-1}^l,$$

$$\mathbf{l}^h = b_0 + b_1\mathbf{p}_1^h + b_2\mathbf{p}_2^h + b_3\mathbf{p}_3^h + \cdots + b_{n-1}\mathbf{p}_{n-1}^h.$$

(5)

For computational reasons, we will use only the first few (three to five) phylogenetic eigenvectors.

When $\mathbf{l}^l$ and $\mathbf{l}^h$ match, the probability of an interaction is high, just like for observed traits $\mathbf{t}^l$ and $\mathbf{t}^h$ above. The values of a_i and b_i have to be found by optimization. Since the values of $\mathbf{l}$ are void of meaning, one cannot assume that the latent vectors align optimally. Therefore, a shift parameter is fitted for the higher trophic level (b_0). Analogous to $\mathbf{T}$, the elements of $\mathbf{L}$, (ℓ_{ij}), are the f_t values computed using equation (3) based on the latent traits: (ℓ_{ij}) = $f_t(l_i^l, l_j^h)$.

Fitting the Model

The tapnet model outlined above can now be fitted to the observed interaction network by adjusting several parameters. As input, the model requires the paired observed traits (zero to many), the phylogeny of each group, the abundance vector for each group, and, for computation of the likelihood, the observed interaction matrix. The model parameters are (i) the width of the trait-matching function (σ in eq. [3]) for each pair of traits, (ii) the width of the trait-matching function for the latent traits, and

(iii) two vectors of parameters for the construction of the latent trait (eq. [5]).

In the optimization using the standard Nelder-Mead algorithm, we used a few tricks to increase the reliability of the model. To ensure identifiability, we constrained a_1 to be positive (i.e., we defined it as e^{a_1}). Otherwise, the exact same values with inverted signs would yield the same fit. Furthermore, any multiple of $\mathbf{a}^l$, $\mathbf{a}^h$ would yield the same prediction. Therefore, we standardized both $\mathbf{l}^l$ and $\mathbf{l}^h$ before entering them into the trait-matching function (eq. [3]).

When fitting multiple networks simultaneously or when using different networks for optimization and prediction, we run into the problem of having different species present in each network. In such cases, we first calculate the eigenvectors of the phylogenetic tree containing all species from all networks. In a second step, we select from these eigenvectors those that are most relevant for the respective network, that is, the eigenvectors most closely correlated with each of the eigenvectors of the tree containing only the species of this particular network. We then fit parameters for only the selected relevant eigenvectors.

Assessing Model Quality

Optimizing the model parameters immediately yields the model's likelihood. Additionally, we may be interested in the latent variables, as they code the (combination of) trait(s) missing in our observed data. For simulated data, we can compare the reconstructed latent traits with those actually simulated. Predicted and observed interactions were additionally compared using other distance measures, such as the Pearson correlation or Bray-Curtis distance, or by summarizing the network structure by means of indices (such as nestedness or specialization). For network indices, we drew 1,000 realizations from the fitted multinomial distribution of each simulated network (see below) and computed network indices for these. Then we computed on which quantile of these 1,000 realization the observed network's index lay (sometimes called the "posterior P value"; Gelman 2005). Ideally, this value

should be .5, indicating no bias in indices in the fitted network.

Simulations

To assess how performance of the tapnet model varies with the characteristics of the data used for fitting, we performed two simulation experiments. The first aimed to evaluate the model's goodness of fit to the interaction network used for estimating parameters, while the second was designed to test the model's accuracy of prediction to a new network. In both experiments we varied six parameters (table 1) using Latin hypercube sampling (McKay et al. 1979). Ranges of numerical parameters were divided into 500 equally spaced intervals, and a single random value was drawn from each interval. For parameters with integer values (e.g., number of observed traits), drawn values were rounded to the nearest integer. In the case of categorical parameters, we randomly sampled 500 times with replacement from the set of possible values. (A prerun with only 100 samples yielded virtually identical results, indicating that 500 runs are sufficient.) Since the number of observed traits varied between zero and four, we drew four sets of 500 values for the width parameter of the trait-matching function. Depending on the number of observed traits of the respective parameter combination, we used only a subset of the four values to calculate matrix $\mathbf{T}$ (or none at all with zero traits).

For each of the 500 parameter combinations, we simulated a data set consisting of relative abundances, phylogenetic trees, and pairs of matching traits. Species abundances were either all set to the same value ($1/m$ or $1/n$, respectively) or drawn from a lognormal distribution with parameters $\mu = 0$ and $\sigma^2 = 1$ and standardized to sum to 1. Trait values were likewise drawn from a lognormal distribution with $\mu = 0$ and $\sigma^2 = 1$. We simulated phylogenetic trees using the function pbtree from the R package phytools (Revell 2012) with a speciation rate of 1 and an extinction rate of 0.

For both experiments, we simulated phylogenies and traits of 30 and 60 lower- and higher-trophic-level species,

Table 1: Parameters varied in the simulation experiments used to assess the model's goodness of fit

Parameter	Type	Range or possible values
Total no. observed interactions (n_{total})	Integer	50–1,000
No. observed traits	Integer	0–4
Type of trait-matching function for observed traits	Categorical	Normal or shifted lognormal
Shape of abundance distribution	Categorical	Uniform or lognormal
Width parameter of trait-matching function for observed traits	Continuous	.05–1
Width parameter of trait-matching function for latent traits	Continuous	.05–2

Note: The width parameter defines the sensitivity of the probability of an interaction to mismatches between traits; small values demand very neat matching for an interaction to be likely.

respectively, for each parameter combination. These data were used to construct a matrix of interaction probabilities **P** according to the tapnet model as described above. While the width parameters of the trait-matching functions for observed and latent traits were systematically varied between simulations, the latent trait linear combination parameters a_i were set to a value of 1 in all simulations. From the interaction probabilities we constructed a network of simulated interactions with a total number of interactions n_{total} by drawing from a multinomial distribution with probabilities $\hat{p}_{i,j}$.

For the first experiment, we then randomly selected 15 (lower-trophic-level) and 30 (higher-trophic-level) species. To these data, we fitted tapnet and assessed the goodness of fit.

For the second experiment, we twice randomly drew a random set of 15 × 30 species. Again, tapnet was fitted to the first but then predicted to the second set. On average, these two sets share half of their species and 25% of their interactions. It thus represents a strong test of predicting to new data. As measures of goodness of fit, we calculated the Bray-Curtis similarity and Spearman rank correlation between the entries of simulated and predicted networks.

Independently Observed versus Network-Derived Abundances

The majority of published interaction networks do not provide independent estimates of the abundances of each species (o_A^{l} and o_A^{h} for lower and higher trophic level, respectively; see eq. [2]). In current network analyses, it is thus customary to use network-derived marginal totals of the network matrix ($\mathbf{O}_{i\cdot}$ and $\mathbf{O}_{\cdot j}$, respectively) as a plug-in instead (e.g., to formulate null model expectations; Vázquez and Aizen 2003; Barber 2007; Blüthgen et al. 2007; Dormann et al. 2009). However, these marginal totals carry the imprint of network structure. At one extreme, pollinators in the region may simply not be attracted by the flowers in the patch under consideration and hence are not present in the network. Or in the case of antagonistic networks, a parasite may reduce the population size of its host to such a degree that interactions are hardly observed, although their intensity is very high (e.g., Barbosa et al. 2017).

We investigated the consequences of using independent versus network-derived abundances on prediction quality with the simulated data. For the same simulated data, we fitted tapnet with the simulated independent abundances once and with the marginal totals of the simulated interaction network once. We then predicted using either the independent abundances of the second simulated network or its marginal totals, respectively. We expect that network-derived abundances will lead to a better model

prediction, simply because they contain information on the structure of the test network.

Case Study

As a demonstration, we use the case study of Tinoco et al. (2017), who compiled data on hummingbird pollination networks in three different habitats (forest, shrubland, and a cattle farm) in the southern Ecuadorian Andes. The data published alongside the article (https://doi.org/10.5061/dryad.j860v) include traits for both plants and hummingbirds as well as external abundance data. These networks are unusually intensively sampled, with 1,288, 3,979, and 2,405 interactions in each of the three habitats, respectively, across 32 plant and 14 hummingbird species, some occurring only in one habitat.

In the case study, we additionally compare the predictive performance of tapnet with three alternative approaches: (1) abundance only; (2) trait-matching and phylogeny GAM, following the ideas of Brousseau et al. (2018); and (3) a similar model using random forest (see the supplemental PDF, available online, for R code and detailed results).

The abundance-only model can be seen as a baseline: it uses only the information on the activity/abundance of the m lower-level and n higher-level species in the validation data v, specifying the estimated probability of interactions: $\hat{\mathbf{P}}_v = (\mathbf{o}_v^{\text{l}}/\sum_{i=1}^{m} o_{v,i}^{\text{l}}) \times (\mathbf{o}_v^{\text{h}}/\sum_{j=1}^{n} o_{v,j}^{\text{h}})$. Multiplying this with the number of observed interactions, $N_{\text{total}} = \sum o^{\text{l}} = \sum o^{\text{h}}$, yields the predicted interaction intensity. Only improving on this model demonstrates the explanatory power of traits and their matching.

Following the approach of Brousseau et al. (2018), we fitted a negative-binomial GAM using 2D splines on the first and second pair of relevant phylogenetic eigenvectors of each group, and we used 1D splines for the observed trait values per species of each group and for the squared difference between traits (representing trait matching). While the original approach used traits and phylogenies to predict binary networks, here we predicted quantitative interaction matrices and additionally used the abundances as predictors. Spline complexity was set to $k = 3$ for univariate and $k = 24$ for 2D splines, and an additional shrinkage was imposed by setting γ to 1.4, both following Brousseau et al. (2018). Note that this approach, as well as the next, implicitly assumes entries of the interaction matrix to be (conditionally) independent (see "Discussion").

The random forest approach was run using the default setting (i.e., 500 trees, trying the rounded-down square root of number of predictors at each split). It was provided with the same information as the previous GAM but using all phylogenetic eigenvectors; it serves as a comparison of the algorithm's flexibility, as it allows for interactions among the predictors.

While tapnet can fit several networks simultaneously (see the supplemental PDF), we employed it akin to the other approaches in a cross-validation setting: the models were fitted to one network and then predicted to the two others in turn. Results were compared using the correlation between predicted and observed interactions.

Models were fit in R using the packages mgcv (Wood 2006) and ranger (Wright and Ziegler 2017); see the supplemental PDF for R code of simulations and the case study.

Results

Simulations

In the first simulation experiment, the correlation between observed and fit networks was overall only moderate (mean value $r = 0.76$ across all 500 parameter combinations). A strong effect of the number of observations was detectable, yielding high correlations between observed and fit networks ($\bar{r} > 0.83$) for networks with more than 0.5 (i.e., 225 interactions in a network with $15 \times 30 = 450$ cells) observations per number of cells (fig. 2, *top left*). Correlation coefficients were lower for networks with equal than for lognormal abundances of all species and higher for the "normal" trait-matching functions than for shifted lognormal. With an increasing number of observed traits, correlation coefficients also increased, indicating the usefulness of observed traits for thinly sampled networks. The two trait-matching function parameters (trait-matching width of latent and observed traits) did not seem to have a clear effect on the correlation between observed and predicted networks. Patterns of variation in Bray-Curtis similarity of observed and predicted networks were similar to those for Pearson's r (results not shown).

In the second simulation experiment, with prediction to a new network, patterns were very similar to those
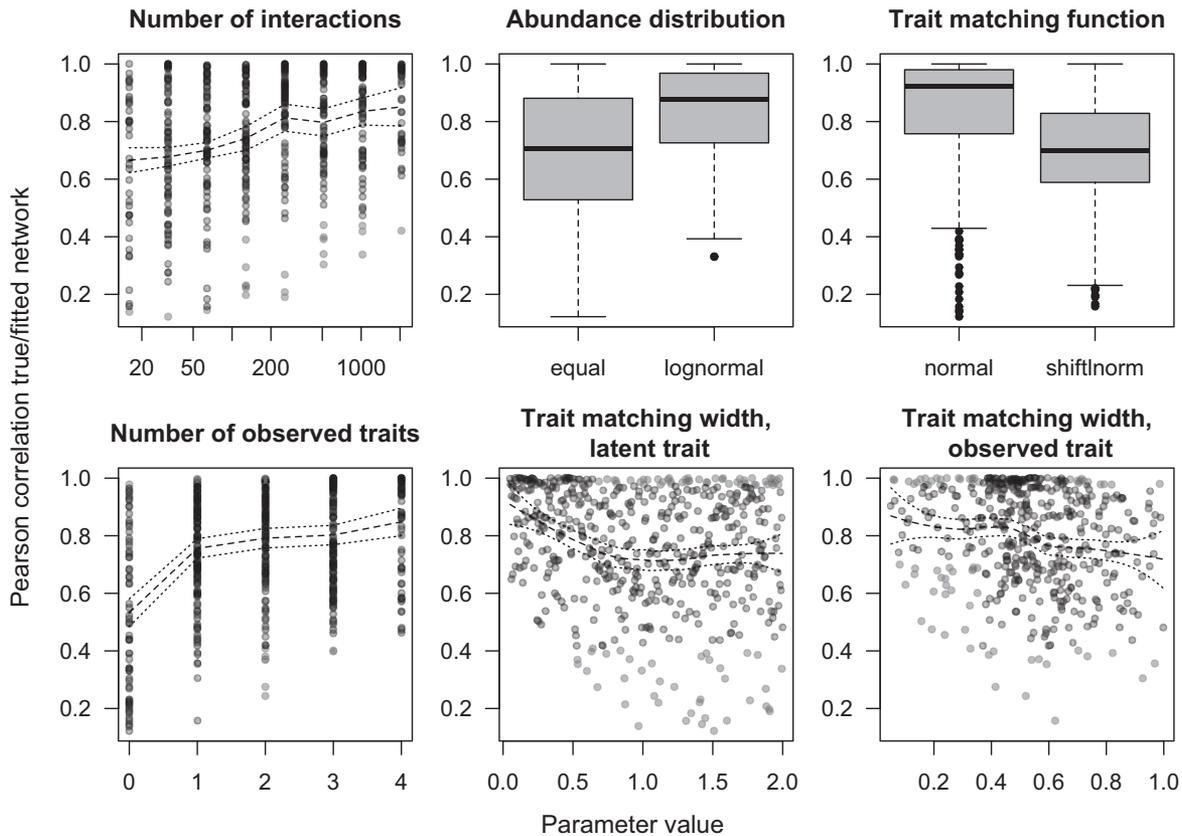

Figure 2: Pearson correlation between fitted model prediction and the simulated interaction network. Scatterplots show the correlation coefficient ρ as a function of six model parameters, whose values were varied simultaneously using Latin hypercube sampling. Data were simulated for 500 parameter combinations. Networks were of size 15×30 species. Lines represent local weighted smoothers and their 95% confidence intervals to indicate trends in the simulations.

in the fitting evaluation, with overall lower correlations ($\bar{r} = 0.42$ and, for densely sampled networks, $\bar{r} = 0.49$). Correlation coefficients depended strongly on the number of observations, abundance distribution, and type of trait-matching function employed (fig. 3). All other parameters had much less effect, and the difference between the two trait-matching functions can thus be seen as two almost separate sets of points in all plots apart from the categorical abundances.

Network indices for random realizations of the fitted network were very similar to those of the observed, usually falling within the 95% confidence interval of the null model (fig. 4). To assess the coverage of the fits and their potential bias, we computed posterior *P* values for each index as mean quantile (see "Assessing Model Quality"). Across the 500 simulations of the first experiment, mean posterior *P* values were .42 for connectance, .44 for NODF (nestedness based on overlap and decreasing fill), .40 for

weighted NODF, and .68 for H_2', all indicating a slightly too generalist estimation of network structure (fig. 4, *bottom row*).

Difference between Independent and Network-Derived Abundances

The fit of tapnet to the simulated data was slightly improved by using the marginal totals rather than the independently "observed" abundances ($\bar{r} = 0.76 \pm 0.22$ [1 SD] for independent abundances compared with $\bar{r} = 0.83 \pm 0.15$ with marginal totals). Also, the prediction to the second simulated network, where half of the species were previously unobserved, improved markedly, from $\bar{r} = 0.41 \pm 0.30$ for independent abundances to $\bar{r} = 0.62 \pm 0.25$ with marginal totals of the new network. Note that using only the marginal totals of the test network
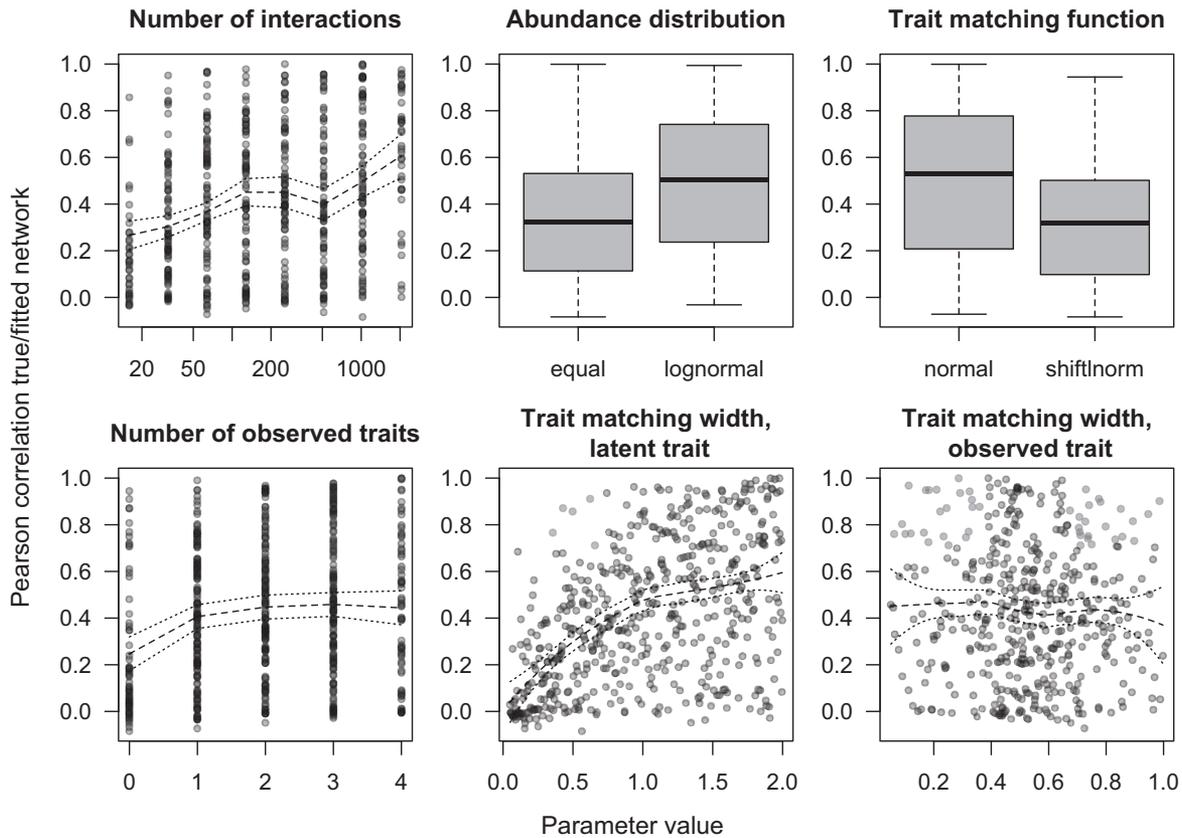

Figure 3: Pearson correlation between model predictions and a new interaction network simulated with identical parameter values. Scatterplots show the correlation coefficient ρ as a function of six model parameters, whose values were varied simultaneously using Latin hypercube sampling. Data were simulated for 500 parameter combinations. Lines represent local weighted smoothers and their 95% confidence intervals to indicate trends in the simulations.
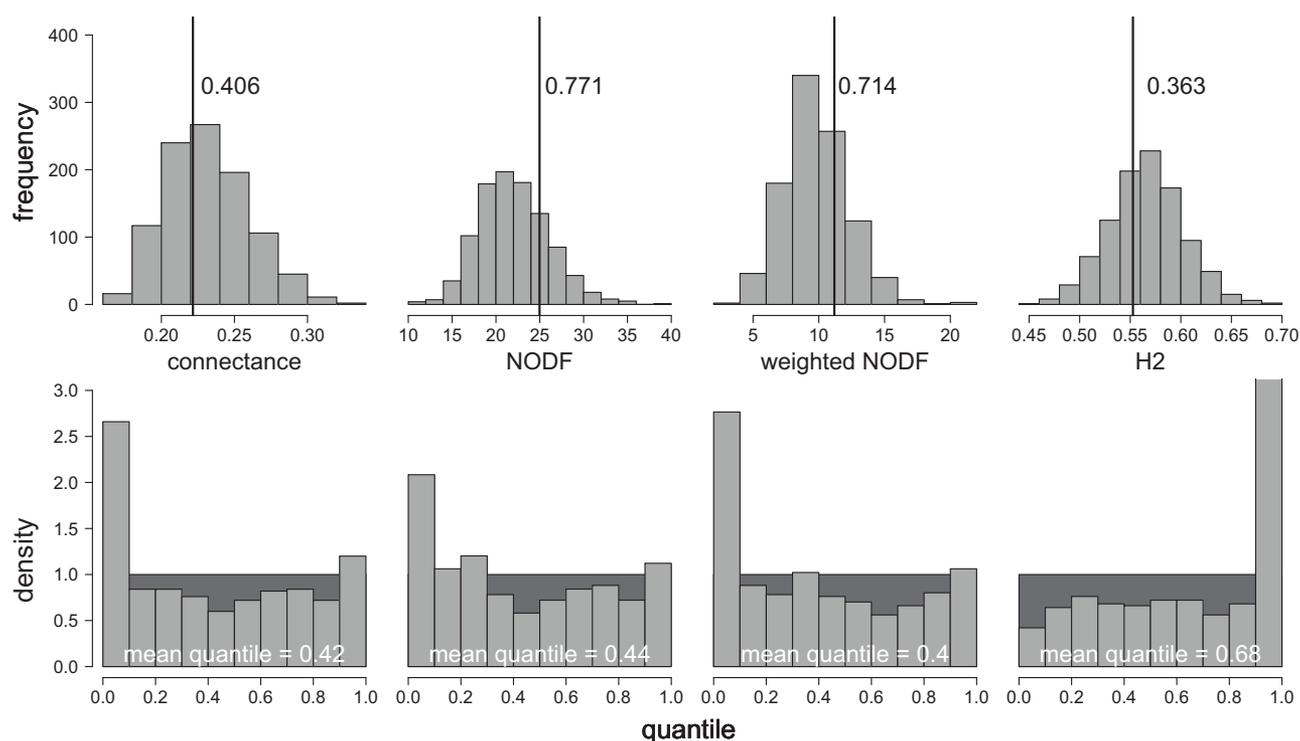
Figure 4: *Top row*, example of network indices for 1,000 realizations of one fitted network, compared with the simulated true value (vertical line; number gives quantile of true value within this frequency distribution). *Bottom row*, quantiles of true network index values across the 500 simulations of experiment 1. Black background indicates ideal uniform distribution. Despite a good overall match, spikes indicate a consistent underestimation of specialization and hence lower connectance, lower nestedness, and higher specialization in the observed network.

(without traits or phylogeny) already predicted interactions well ($\bar{r} = 0.66 \pm 0.17$).

Case Study

The tapnet model fitted the three hummingbird-flower networks better than the abundance-only model and the GAM but worse than the random forest approach (table 2). On cross validation, when fitting to one and predicting to the other two networks in turn, the tapnet approach was no better or worse than the abundance-only

approach and the random forest approach but was substantially better than the GAM (table 3). Interestingly, all approaches except the trait-neutral, abundance-only approach declined dramatically in performance from the training to the test data. The GAM, for example, fitted the data moderately but held no predictive power for the test data. Random forest, although reporting an extremely good fit, decreased to the level of tapnet and the abundance-only approach on the test data. This drop in performance from fitting to predicting suggests that all statistical approaches overfitted, sometimes heavily.

Discussion

In recent years, two main but not mutually exclusive lines of modeling approaches to predict network structure have emerged (as reviewed in Valdovinos 2019). Neutral models assume all species to be similar and generalist and hence describe an expectation for network structure primarily based on sampling intensity and abundance distributions. Indeed, such approaches are often used as a null model against which to gauge the effect of interaction preferences (e.g., Vázquez and Aizen 2003; Blüthgen et al. 2006; Dormann et al. 2009). In contrast, interaction constraint

Table 2: Pearson correlation coefficients between fitted and observed hummingbird-flower network (Tinoco et al. 2017) for the four approaches and their means across the three habitats

Approach	Forest	Shrub	Farm	Mean
Abundance only	.25	.09	.49	.28
Random forest	.93	.93	.91	.92
Tapnet	.57	.60	.65	.61
Generalized additive model	.56	.29	.40	.42

Note: Approaches are sorted by ranking in the cross-validation performance (table 3).

Table 3: Pearson correlation of cross validation by predicting with a model fitted to one habitat to the other habitats (indicated by a right arrow)

Approach	F→S	F→C	S→F	S→C	C→F	C→S	Mean
Abundance only	.09	.49	.25	.49	.25	.09	.28
Random forest	.23	.36	.33	.14	.40	.17	.27
Tapnet	.21	.46	.12	.53	.15	.11	.26
Generalized additive model	−.01	.33	.26	.13	−.01	−.02	.11

Note: Habitats are forest (F), shrubland (S), and a cattle farm (C). For cross-validation log likelihoods, which show the same result, see the supplemental PDF.

models focus on why some links are not present, typically using species traits and their match across the groups as explanatory features (Santamaría and Rodríguez-Gironés 2007; Bartomeus et al. 2016). In the approach presented here, we combine both by allowing abundances to contribute or even entirely dominate the prediction but use matching between observed traits and between latent traits as constraints.

The results so far are both promising and sobering. On the one hand, we demonstrated that we can fit the observed data well with the tapnet approach and reasonably predict simulated data. We can also predict with moderate accuracy networks from other habitats, and the tapnet approach did that better than some previous attempts to combine observed and latent traits (Brousseau et al. 2018). On the other hand, by far the most important predictor for our case study was species abundance, which makes it easy for any approach and prevents the trait-matching strength of tapnet to play out. Thus, a neutral model, using only abundances, was as good at prediction (but not fit)—and at the same time much simpler—than tapnet or random forest. We think that at least two factors contribute to this finding: (1) the hummingbird-flower network is not very specialized, despite featuring the most spectacular sword-billed hummingbird (*Ensifera ensifera*), whose pollination-adapted bill is longer than its body; and (2) when abundance is very important, the log likelihood becomes very shallow and difficult to improve on. Also, abundance may be correlated with a trait relevant for interactions, although in this case there was no correlation between bill/corolla size and abundances ($r = 0.019$ and $−0.15$, respectively; see the supplemental PDF).

In simulations with many (>500) observations, tapnet predictions were very reasonable, while in the case study all approaches fared relatively poorly. Our simulations included processes deemed to be most important for determining network structure (compared with the list in Valdovinos 2019). This suggests either that mutualist networks may simply be extremely noisy and under low evolutionary pressure or that the current trait-matching concepts are not good enough for describing, across networks, the processes that drive interactions. This suggests a strong context dependence of interactions, depending, probably

substantially, on intraspecific trait variation (Laughlin et al. 2012), behavioral complexity (Kaiser-Bunbury et al. 2010; Morán-López et al. 2020), competition within guilds (Vandermeer 2004; Saavedra et al. 2013), nonlinear frequency dependence (Benadi and Pauw 2018), and environmental conditions more generally (for a review, see Valdovinos 2019). Future research across many different networks has to show whether abundance is consistently such an important predictor for interaction frequencies.

The Chicken-and-Egg Problem of Abundances and Network Structure

Not all studies record independent abundances—for example, by estimating floral cover, sweep-netting insects, or the like—and this is more common for birds than for insects. Without such independent data for abundance, its role for determining interaction frequencies cannot be determined. One frequent "solution" is to use the observed interactions of each species (the marginal totals of the interaction matrix) as a surrogate for its abundance. This approach has been rightly criticized as confounding the effect of abundances on network interactions with the effect of network structure on abundances—that is, it is the chicken-and-egg problem of network interactions (Fort et al. 2016; Dormann et al. 2017). This conflict was also detectable in our simulations, where the tapnet prediction to the test data was substantially improved by using the test network's marginal total as the predictor. Clearly, these surrogate abundances carry some information, beyond abundance, on network structure and hence interaction intensity.

Flower visitation networks are, in general, only moderately "ecologically specialized" (sensu Armbruster 2017; see, e.g., Blüthgen et al. 2007; Schleuning et al. 2012; Zanata et al. 2017), suggesting that neither plants nor pollinators depend crucially on a specific (set of) species to interact with. As a consequence, network structure and species abundances are strongly linked. In these cases, marginal totals may arguably be used in lieu of independent abundances. Our case study of a plant-pollinator system shows that this is not the correct approach. Indeed, when predicting the fitted model once with the external and once with the marginal abundances, we find a dramatically

better performance for the marginal abundances (across the three habitats, the predictive correlation is $\bar{r} = 0.75$, compared with a meager $\bar{r} = 0.26$ with independent abundances; see the supplemental PDF). Clearly, marginal abundances contain an a priori unknown amount of information, representing the outcome of within-guild interactions, variable activity of individuals, selection cues (scent, visual signals), and so forth.

How Independent Are Observed Interactions?

We believe that treating observed interactions as independent data is statistically incorrect (see the introduction), thus questioning the likelihood used by Rohr et al. (2016) and Brousseau et al. (2018) as well as the implicit independence assumption in Pichler et al. (2020). Therefore, we here used a (network-wide) multinomial distribution, as had been suggested by Vázquez et al. (2009), accommodating the compositional nature of the data. While it recognizes the nonindependence of observations, it does not thereby automatically capture the processes behind it. That means that while the inference based on this distribution is probably correct, our multinomial approach may not result in better predictions until the drivers of nonindependence are represented in the model (e.g., intraguild competition). Additionally, this approach has two statistically relevant implications. First, we get only a single likelihood value for a network rather than *nm*. Second, as a consequence we assess the fit of the entire model, without any chance of adapting only the fits of some specific species or interaction, as is the case for the GAM and random forest approaches used here for comparison. While we regard it as a more correct representation of the data, it also severely limits the type of statistical approaches that can be used for predicting network interactions.

Olito and Fox (2015) focus on a comparison of network indices produced by their predictive approach with those of the observed network. They conclude that even similar networks may have rather different index values (and different networks similar indices), making such indices a poor target for optimization. In our simulated networks, we were able to fit networks so that their index values centered on the observed value. This illustrates that tapnet did manage to fit network indices in line with the observations as a by-product of the multinomial likelihood in principle.

Traits, Observed and Latent

Following the lead of previous studies (in particular, Rohr et al. 2016), tapnet uses phylogenetic information to construct (not necessarily phylogenetically conserved) latent traits to improve fit to data. Clearly, these are only a sta-

tistical placeholder for actual but unobserved ecological traits. Beyond the obvious but difficult-to-measure sensory interaction cues (Junker et al. 2013), traits related to optimal foraging should also be considered here—from both groups of interacting species. As Pyke (2016) exemplifies, the fitness benefit of pollination for plants depends on the pollen-transfer efficiency of pollinators, and too-high nectar rewards may give an incentive to inefficient visitors. Latent traits may thus reflect a complex and fine-tuned pair of matching sets of traits, without obvious interpretation.

A corollary of the possibility of fitting trait pairs is a risk of identifying spurious characteristics of species (Mlambo 2014). At present, interaction traits are almost exclusively morphological (for an exception, see Junker et al. 2013), but phenology can be incorporated similarly. In the future, both ecophysiological and genetic sampling may become sensitive enough to extend research into interaction traits, for example, related to vision in the ultraviolet (e.g., Rae and Vamosi 2013), scent (e.g., Wright and Schiestl 2009), or ultrasonic sound (e.g., Simon et al. 2019). Until such devices become available, latent traits are a statistical stand-in for what really makes species interact.

Conclusion

The approach we have presented here predicts network interactions for new networks, conditional on trait, abundance/activity, and phylogenetic data for the new network. It is flexible enough to include any type of function translating trait matches into interaction probabilities. As a side effect, it quantifies the importance of abundance relative to traits for network interactions. Future applications will have to assess the importance of traits across different types of networks, testing the assumption of many network studies that traits are the driving force of network interactions.

Statement of Authorship

Conceptualization: G.B., C.F.D., J.F., D.P.V.; methods and software development: G.B., C.F.D., R.S.; code review: G.B., C.F.D., J.F., R.S.; simulations and analyses:

G.B., C.F.D.; writing of original draft: G.B., C.F.D.; reviewing and editing: G.B., C.F.D., J.F., R.S., D.P.V.

Data and Code Availability

All data and code are available as part of the tapnet package on GitHub (https://github.com/biometry/tapnet) and CRAN (https://cran.r-project.org).

Literature Cited

Allesina, S. 2011. Predicting trophic relations in ecological networks: a test of the allometric diet breadth model. Journal of Theoretical Biology 279:161–168.

Armbruster, W. S. 2017. The specialization continuum in pollination systems: diversity of concepts and implications for ecology, evolution and conservation. Functional Ecology 31:88–100.

Barber, M. 2007. Modularity and community detection in bipartite networks. Physical Review E 76:066102.

Barbosa, M., G. W. Fernandes, O. T. Lewis, and R. J. Morris. 2017. Experimentally reducing species abundance indirectly affects food web structure and robustness. Journal of Animal Ecology 86:327–336.

Bartomeus, I., D. Gravel, J. M. Tylianakis, M. A. Aizen, I. A. Dickie, and M. Bernard-Verdier. 2016. A common framework for identifying linkage rules across different types of interactions. Functional Ecology 30:1894–1903.

Bascompte, J., and P. Jordano. 2014. Mutualistic networks. Princeton University Press, Princeton, NJ.

Beckerman, A. P., O. L. Petchey, and P. H. Warren. 2006. Foraging biology predicts food web complexity. Proceedings of the National Academy of Sciences of the USA 103:13745–13749.

Benadi, G., and A. Pauw. 2018. Frequency dependence of pollinator visitation rates suggests that pollination niches can allow plant species coexistence. Journal of Ecology 106:1892–1901.

Blüthgen, N., F. Menzel, and N. Blüthgen. 2006. Measuring specialization in species interaction networks. BMC Ecology 6:9.

Blüthgen, N., F. Menzel, T. Hovestadt, B. Fiala, and N. Blüthgen. 2007. Specialization, constraints, and conflicting interests in mutualistic networks. Current Biology 17:341–346.

Brousseau, P.-M., D. Gravel, and I. T. Handa. 2018. Trait matching and phylogeny as predictors of predator-prey interactions involving ground beetles. Functional Ecology 32:192–202.

Chittka, L., and N. E. Raine. 2006. Recognition of flowers by pollinators. Current Opinion in Plant Biology 9:428–435.

Clarke, K. R., and R. M. Warwick. 1999. The taxonomic distinctness measure of biodiversity: weighting of step lengths between hierarchical levels. Marine Ecology Progress Series 184:21–29.

Crea, C., R. A. Ali, and R. Rader. 2016. A new model for ecological networks using species level traits. Methods in Ecology and Evolution 7:232–241.

Dehling, D. M., P. Jordano, H. M. Schaefer, K. Böhning-Gaese, and M. Schleuning. 2016. Morphology predicts species' functional roles and their degree of specialization in plant-frugivore interactions. Proceedings of the Royal Society B 283:20152444.

Desjardins-Proulx, P., I. Laigle, T. Poisot, and D. Gravel. 2017. Ecological interactions and the Netflix problem. PeerJ 5:e3644.

Dormann, C. F., and N. Blüthgen. 2017. Food webs versus interaction networks: principles, pitfalls and perspectives. Pages 1–11 in J. C. Moore, P. C. de Ruiter, K. S. McCann, and V. Wolters, eds. Adaptive food webs: stability and transitions of real and model ecosystems. Cambridge University Press, Cambridge.

Dormann, C. F., N. Blüthgen, J. Fründ, and B. Gruber. 2009. Indices, graphs and null models: analyzing bipartite ecological networks. Open Ecology Journal 2:7–24.

Dormann, C. F., J. Fründ, and H. M. Schaefer. 2017. Identifying causes of patterns in ecological networks: opportunities and limitations. Annual Reviews in Ecology, Evolution, and Systematics 48:559–584.

Fort, H., D. P. Vázquez, and B. L. Lan. 2016. Abundance and generalization in mutualistic networks: solving the chicken-and-egg dilemma. Ecology Letters 19:4–11.

Gelman, A. 2005. Comment: fuzzy and Bayesian p-values and u-values. Statistical Science 20:380–381.

Gravel, D., T. Poisot, C. Albouy, L. Velez, and D. Mouillot. 2013. Inferring food web structure from predator-prey body size relationships. Methods in Ecology and Evolution 4:1083–1090.

Guénard, G., P. Legendre, and P. Peres-Neto. 2013. Phylogenetic eigenvector maps: a frame-work to model and predict species traits. Methods in Ecology and Evolution 4:1120–1131.

Guimarães, P. R., Jr., P. Jordano, and J. N. Thompson. 2011. Evolution and coevolution in mutualistic networks. Ecology Letters 14:877–885.

Houlahan, J. E., S. T. Mckinney, T. M. Anderson, and B. J. McGill. 2017. The priority of prediction in ecological understanding. Oikos 126:1–7.

Ives, A. R., and H. C. J. Godfray. 2006. Phylogenetic analysis of trophic associations. American Naturalist 168:E1–E14.

Junker, R., N. Blüthgen, T. Brehm, J. Binkenstein, J. Paulus, H. M. Schaefer, and M. Stang. 2013. Specialization on traits as basis for the niche-breadth of flower visitors and as structuring mechanism of ecological networks. Functional Ecology 27:329–341.

Junker, R., N. Höcherl, and N. Blüthgen. 2010. Responses to olfactory signals reflect network structure of flower-visitor interactions. Journal of Animal Ecology 79:818–823.

Kaiser-Bunbury, C. N., S. Muff, J. Memmott, C. B. Müller, and A. Caflisch. 2010. The robustness of pollination networks to the loss of species and interactions: a quantitative approach incorporating pollinator behaviour. Ecology Letters 13:442–452.

Laughlin, D. C., C. Joshi, P. M. van Bodegom, Z. A. Bastow, and P. Z. Fulé. 2012. A predictive model of community assembly that incorporates intraspecific trait variation. Ecology Letters 15:1291–1299.

Loeuille, N., and M. Loreau. 2005. Evolutionary emergence of size-structured food webs. Proceedings of the National Academy of Sciences of the USA 102:5761–5766.

McKay, M., R. Beckman, and W. Conover. 1979. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 21:239–245.

Mlambo, M. C. 2014. Not all traits are "functional": insights from taxonomy and biodiversity-ecosystem functioning research. Biodiversity and Conservation 23:781–790.

Moore, J. C., P. C. de Ruiter, K. S. McCann, and V. Wolters. 2017. Adaptive food webs: stability and transitions of real and model ecosystems. Cambridge University Press, Cambridge.

Morales-Castilla, I., M. G. Matias, D. Gravel, and M. B. Araújo. 2015. Inferring biotic interactions from proxies. Trends in Ecology and Evolution 30:347–356.

Morán-López, T., W. D. Espíndola, B. S. Vizzachero, A. Fontanella, L. Salinas, C. Arana, G. Amico, M. A. Pizo, T. A. Carlo, and J. M.

Morales. 2020. Can network metrics predict vulnerability and species roles in bird-dispersed plant communities? not without behaviour. Ecology Letters 23:348–358.

Muñoz, M. C., H. M. Schaefer, K. Böhning-Gaese, and M. Schleuning. 2017. Importance of animal and plant traits for fruit removal and seedling recruitment in a tropical forest. Oikos 126:823–832.

Olito, C., and J. W. Fox. 2015. Species traits and abundances predict metrics of plant-pollinator network structure, but not pairwise interactions. Oikos 124:428–436.

Pearse, I. S., and F. Altermatt. 2013. Predicting novel trophic interactions in a non-native world. Ecology Letters 16:1088–1094.

Pearse, I. S., D. J. Harris, R. Karban, and A. Sih. 2013. Predicting novel herbivore-plant interactions. Oikos 122:1554–1564.

Peralta, G. 2016. Merging evolutionary history into species interaction networks. Functional Ecology 30:1917–1925.

Petchey, O. L., A. P. Beckerman, J. O. Riede, and P. H. Warren. 2008. Size, foraging, and food web structure. Proceedings of the National Academy of Sciences of the USA 105:4191–4196.

Pichler, M., V. Boreux, A. Klein, M. Schleuning, and F. Hartig. 2020. Machine learning algorithms to infer trait-matching and predict species interactions in ecological networks. Methods in Ecology and Evolution 11:281–293.

Poisot, T., A. R. Cirtwill, K. Cazelles, D. Gravel, M.-J. Fortin, and D. B. Stouffer. 2016. The structure of probabilistic networks. Methods in Ecology and Evolution 7:303–312.

Pomeranz, J. P. F., R. M. Thompson, T. Poisot, and J. S. Harding. 2019. Inferring predator-prey interactions in food webs. Methods in Ecology and Evolution 10:356–367.

Pyke, G. H. 1984. Optimal foraging theory: a critical review. Annual Review of Ecology and Systematics 15:523–575.

———. 2016. Plant-pollinator co-evolution: it's time to reconnect with optimal foraging theory and evolutionarily stable strategies. Perspectives in Plant Ecology, Evolution and Systematics 19:70–76.

Rae, J. M., and J. C. Vamosi. 2013. Ultraviolet reflectance mediates pollinator visitation in *Mimulus guttatus*. Plant Species Biology 28:177–184.

Revell, L. J. 2012. Phytools: an R package for phylogenetic comparative biology (and other things). Methods in Ecology and Evolution 3:217–223.

Rohr, R. P., R. E. Naisbit, C. Mazza, and L.-F. Bersier. 2016. Matching-centrality decomposition and the forecasting of new links in networks. Proceedings of the Royal Society B 283: 20152702.

Rooney, N., K. McCann, G. Gellner, and J. C. Moore. 2006. Structural asymmetry and the stability of diverse food webs. Nature 442:265–269.

Saavedra, S., R. P. Rohr, V. Dakos, and J. Bascompte. 2013. Estimating the tolerance of species to the effects of global environmental change. Nature Communications 4:2350.

Santamaría, L., and M. A. Rodríguez-Gironés. 2007. Linkage rules for plant-pollinator networks: trait complementarity or exploitation barriers? PLoS Biology 5:e31.

Schleuning, M., J. Fründ, A.-M. Klein, S. Abrahamczyk, R. Alarcón, M. Albrecht, G. K. S. Andersson, et al. 2012. Specialization of mutualistic interaction networks decreases toward tropical latitudes. Current Biology 22:1925–1931.

Sebastián-González, E., M. Moleón, J. P. Gibert, F. Botella, P. Mateo-Tomás, P. P. Olea, P. R. Guimarães, and J. A. Sánchez-

Zapata. 2016. Nested species-rich networks of scavenging vertebrates support high levels of interspecific competition. Ecology 97:95–105.

Simon, R., F. Matt, V. Santillan, M. Tschapka, M. Tuttle, and W. Halfwerk. 2019. An ultrasound absorbing inflorescence zone enhances echo-acoustic contrast of bat-pollinated cactus flowers. bioRxiv, https://doi.org/10.1101/2019.12.28.890046.

Spiesman, B. J., and C. Gratton. 2016. Flexible foraging shapes the topology of plant-pollinator interaction networks. Ecology 97: 1431–1441.

Stephens, D. W., and J. R. Krebs. 1986. Foraging theory. Monographs in Behavior and Ecology. Princeton University Press, Princeton, NJ.

Tinoco, B. A., C. H. Graham, J. M. Aguilar, and M. Schleuning. 2017. Effects of hummingbird morphology on specialization in pollination networks vary with resource availability. Oikos 126: 52–60.

Valdovinos, F. S. 2019. Mutualistic networks: moving closer to a predictive theory. Ecology Letters 22:1517–1534.

Valdovinos, F. S., E. L. Berlow, P. M. de Espanés, R. Ramos-Jiliberto, D. P. Vázquez, and N. D. Martinez. 2018. Species traits and network structure predict the success and impacts of pollinator invasions. Nature Communications 9:2153.

Vandermeer, J. 2004. Coupled oscillations in food webs: balancing competition and mutualism in simple ecological models. American Naturalist 163:857–867.

Vázquez, D. P., and M. A. Aizen. 2003. Null model analyses of specialization in plant-pollinator interactions. Ecology 84:2493–2501.

Vázquez, D. P., N. Chacoff, and L. Cagnolo. 2009. Evaluating multiple determinants of the structure of plant-animal mutualistic networks. Ecology 90:2039–2046.

Veits, M., I. Khait, U. Obolski, E. Zinger, A. Boonman, A. Goldshtein, K. Saban, et al. 2019. Flowers respond to pollinator sound within minutes by increasing nectar sugar concentration. Ecology Letters 22:1483–1492.

Vizentin-Bugoni, J., V. J. Debastiani, V. A. G. Bastazini, P. K. Maruyama, and J. H. Sperry. 2020. Including rewiring in the estimation of the robustness of mutualistic networks. Methods in Ecology and Evolution 11:106–116.

Warton, D. I., F. G. Blanchet, R. B. O. Hara, O. Ovaskainen, S. Taskinen, S. C. Walker, and F. K. C Hui. 2015. So many variables: joint modeling in community ecology. Trends in Ecology and Evolution 30:766–779.

Wood, S. N. 2006. Generalized additive models: an introduction with R. Chapman & Hall/CRC, London.

Wright, G. A., and F. P. Schiestl. 2009. The evolution of floral scent: the influence of olfactory learning by insect pollinators on the honest signalling of floral rewards. Functional Ecology 23:841–851.

Wright, M. N., and A. Ziegler. 2017. ranger: a fast implementation of random forests for high dimensional data in C++ and R. Journal of Statistical Software 77:1–17.

Zanata, T. B., B. Dalsgaard, F. C. Passos, P. A. Cotton, J. J. Roper, P. K. Maruyama, E. Fischer, et al. 2017. Global patterns of interaction specialization in bird-flower networks. Journal of Biogeography 44:1891–1910.