ECOGRAPHY

Research article

Layer-specific imprints of traits within a plant–herbivore–predator network – complementary insights from complementary methods

Kate L. Wootton[1,2], **F. Guillaume Blanchet**[3,7,8], **Andrew Liston**[4], **Tommi Nyman**[5], **Laura G. A. Riggi**[1], **Jens-Peter Kopelke**[6], **Tomas Roslin**[1] and **Dominique Gravel**[3]

[1]Department of Ecology, Swedish University of Agricultural Sciences, Uppsala, Sweden
[2]School of Biological Sciences, University of Canterbury, Christchurch, New Zealand
[3]Département de biologie, Université de Sherbrooke, Sherbrooke, Québec, Canada
[4]Senckenberg German Entomological Institute Muncheberg, Brandenburg, Germany
[5]Department of Ecosystems in the Barents Region, Norwegian Institute of Bioeconomy Research, Svanvik, Norway
[6]Senckenberg Gesellschaft fur Naturforschung Frankfurt am Main, Hessen, Germany
[7]Département de mathématiques, Université de Sherbrooke, Sherbrooke, Québec, Canada
[8]Département des sciences de la santé communautaire, Université de Sherbrooke, Sherbrooke, Québec, Canada

Correspondence: Kate L. Wootton (kate.l.wootton@gmail.com)

Ecography
2024: e07028
doi: 10.1111/ecog.07028

Subject Editor: Jose M. Montoya
Editor-in-Chief: Miguel Araújo
Accepted 29 November 2023

Who interacts with whom is a key question in community and network ecology. The concept that these interactions may be driven by a match between the traits of consumer and resource species is known as trait-matching. If trait-matching would allow for general predictions of interaction structure based on sufficiently few and easily-measurable traits, then this approach could replace the laborious description of each individual pairwise interaction. To resolve imprints of trait-matching in a species-rich tri-trophic *Salix*–galler–parasitoid network, and to identify the most relevant traits, we applied five different methods, each approaching the same phenomenon from a different perspective. As traits, we used, body sizes, gall type (position on plant, structure of gall) and phenology, among others, as well as phylogenetic proxies. When jointly applied, the methods demonstrate distinctly different imprints of traits within the two bipartite network elements (*Salix*–galler versus galler–parasitoid interactions). Of the galler–parasitoid sub-network's interactions, approximately half were explainable by the species traits used; of the *Salix*–galler sub-network's interactions, traits explained at most two-fifths. Gall type appeared to be the most important structuring trait in both networks. Phylogeny explained as much, or more than did our tested traits, suggesting that traits may be conserved and phylogeny therefore an effective proxy. Overall, the more specialized structure of the *Salix*–galler network versus the more nested structure of the galler–parasitoid network meant that different methods were more effective at capturing interactions and interaction structure in the different sub-networks. Thus, our analysis reveals how structuring impacts may vary even between levels within the same multitrophic network, and calls for comparative analyses of trait matching across a wide set of systems and methods.

Keywords: fourth-corner, K-nearest neighbour, matching-centrality, parasitoid, Random forest, *Salix*, willow-galling sawflies

NORDIC SOCIETY OIKOS

www.ecography.org

Introduction

Traits are morphological, behavioural, or chemical characteristics of organisms that differ among individuals and/or groups (Violle et al. 2007) and that restrict whom individuals can effectively interact with – such as whom they can eat or be eaten by in a feeding interaction. There is accumulating evidence that traits play a structuring role in ecological networks such as food webs, and increasing interest in the use of traits to predict interactions and network structure (Morales-Castilla et al. 2015, Bartomeus et al. 2016, Laigle et al. 2018, Pichler et al. 2020). Traits may provide the key to move from idiosyncratic observations of pairwise interactions, to infer general mechanisms underlying interactions and dynamics (Bartomeus et al. 2016).

Encouragingly, the evidence now suggests that relatively few traits may be needed to explain interactions among large sets of species (Eklöf et al. 2013, Brousseau et al. 2018). A number of methods using species' traits such as body size have been able to replicate both food-web structures and dynamics (Williams and Martinez 2000, Schneider et al. 2012, Gravel et al. 2016, Rohr et al. 2016). Yet, despite these theoretical and methodological developments, we still know relatively little both about which traits structure interaction networks and how they do so.

One of the most topical concepts of how traits structure networks is that of trait matching (Rossberg 2008, Dehling et al. 2014, Gravel et al. 2016, Brousseau et al. 2018). Trait matching holds that the set of possible interactions is constrained by a match between consumers' foraging traits and resources' vulnerability traits. Based on this concept, similar species should share similar interactions (Bartomeus et al. 2016, Desjardins-Proulx et al. 2017). Models using trait matching have developed, with progressively better results, from the cascade model (Cohen and Newman 1985), where species along a trait continuum feed on those below them, to the niche model, which uses the same continuum but yields a narrower match between consumer and resource traits by restricting each species to a foraging range (Williams and Martinez 2000). Further developments include the probabilistic niche model (Williams et al. 2010), which replaced the feeding range with a probability distribution where interaction probabilities decrease away from an optimal point (Williams and Purves 2011), and the group model (Allesina and Pascual 2009), which finds groups of species that share the same interactions. Body size is used most commonly for the trait continuum in the hierarchical models (cascade and niche models) and phylogeny for the group model (Staniczenko et al. 2014), although any trait could potentially be used. Our understanding of trait-based interactions is continuously being expanded by new ideas and models (Gravel et al. 2013, Bartomeus et al. 2016, Laigle et al. 2018).

In addition to trait matching between predator and prey, trait distribution within a community also structures networks (Bartomeus et al. 2016, Laigle et al. 2018). For example, generalist consumers – those that target many prey species – tend to display wider intraspecific phenotypic variation than specialists which allows them, as a species, to target a wider range of prey types (Gibert and DeLong 2017). This can give rise to a nested structure, particularly in bipartite networks. In nested networks, specialists interact with subsets of the species generalists interact with (Bascompte et al. 2003). However, this pattern primarily emerges at a binary level and, when interaction strengths are taken into account, binary nested networks often turn out to be quantitatively unnested, meaning that species' preferences weaken competition (Staniczenko et al. 2013). The matching-centrality framework (Rohr et al. 2010, 2016), which we explore in greater detail in the methods, aims to capture the importance of both of these aspects. In summary, the match between a predator's foraging traits and prey's vulnerability traits constrains the likelihood of an interaction between two given species (Boukal 2014), whereas overall network structure also depends on the distribution of all traits within the species community (Bartomeus et al. 2016, Gravel et al. 2016, Laigle et al. 2018).

Different models of how networks are organised by traits will call for different methods for extracting the corresponding signal from empirical data. To date, trait-based analyses of network structure have generally been approached using regression-based methods (Gravel et al. 2013, Pearse and Altermatt 2013, Pomeranz et al. 2019). However, the considerations above suggest that relationships between traits and interactions may be more complex and require more flexible models. Methods such as Random forest (Breiman 2001) may then be more appropriate for capturing relationships between traits and network structure (Desjardins-Proulx et al. 2017, Laigle et al. 2018, Pichler et al. 2020).

Finally, there is the challenge of what traits to measure. While the likely mechanics of the interaction may guide us in selecting what traits to measure (Wootton et al. 2023), there is always the risk that a key trait may remain unmeasured. Since related species tend to share similar traits, phylogenetic distance may serve as a proxy of morphological, chemical and behavioural similarity, and has often been found effective at explaining network structure either instead of or in complement to traits measured, such as body size (Naisbit et al. 2012, Eklöf et al. 2013, Rohr and Bascompte 2014, Gray et al. 2015).

In this study, we use a uniquely extensive data set consisting of *Salix* species, *Salix*–galling sawflies, and their parasitoids (Kopelke et al. 2017), to explore the role of trait-matching in an unusually well-characterized network. By applying five different methods to the same quest, we dissect the relationship between species' traits and network structure from multiple dimensions – adding phylogeny to approximate unmeasured traits. The methods we selected are K-nearest neighbour (KNN) (Desjardins-Proulx et al. 2017); indirect matching-centrality (IMC) (Rohr et al. 2016); Random forest; generalized linear models (GLM), and the Fourth corner approach. By combining insights from these methods, we can ask questions about how much traits or phylogeny can actually tell us (by comparing methods using traits or phylogeny to equivalent methods that do not), how the network is structured, and overall get a deeper understanding about the ecology of the network. In particular, we used KNN to ask whether similar species are more likely to share interactions, and if this is due to their traits and/or phylogeny. We ask whether parametric approaches (IMC, GLM, Fourth corner) – where the functional response between traits and interactions is predefined (Lucas 2020) – adequately represent the relationship between traits and network structure. As an alternative, we address whether more flexible non-parametric methods (KNN and Random forest) – where the relationship is learned algorithmically – are more appropriate. Finally, we ask which trait(s) are most important in structuring networks, across all methods, and delve into the ecology of the focal species to explore why.

Material and methods

Empirical network data

We focus on the data set presented by Kopelke et al. (2017). After pruning of rare species (Supporting information), it contains 115 unique *Salix*–galler interactions (out of a possible 3080) between 35 willow species (*Salix*) and their hybrids and 88 willow-galling sawfly species (Hymenoptera: Tenthredinidae: Nematinae: Euurina), and 812 interactions (out of a possible 4488) between the 88 gallers and 51 of their natural enemies. While these natural enemies span from hymenopteran true parasitoids to coleopteran, lepidopteran, dipteran, and hymenopteran inquilines, for simplicity we will henceforth refer to the third trophic layer as 'parasitoids'. The original data were collected during 641 visits at 374 sites across Europe over 29 years and included a total of 165,424 individual galls sampled. Interactions were identified by collecting galls from individual willows (identifying a *Salix*–galler interaction) and then rearing galls in the lab until either the adult galler or a parasitoid emerged (identifying a galler–parasitoid interaction) (Kopelke 1999, Kopelke et al. 2003). The data from each site and visit were combined into one aggregated interaction matrix, where we classified species pairs into three categories: those that interacted (1), those that did not (0), and those which were never seen to co-occur at the same site.

Trait data

Species-level data

For a subset of *Salix* and galler species, we measured traits in the field at sites ranging from Greece to northern Norway (Supporting information). Other traits were obtained by consulting the literature and experts in the field. For exact details on trait measurements and collection, we refer the reader to the Supporting information.

Salix traits included characteristics of the leaves (specific leaf area (SLA), toughness, size, hairiness and thickness) and trees (height). For those species not found in the field, we consulted experts and the literature (Jonsell and Karlsson 2000). *Salix* phenology were obtained through expert consultation. The level of glucoside defense we ascertained from the literature (Julkunen-Tiitto 1989, Meier et al. 1992). For the latter, we used a coarse measure of total glucoside concentration, as data on specific glucoside compounds proved insufficient (Supporting information).

Galler traits included characteristics of the gall itself (type, volume, wall thickness, position on the plant), and physical and behavioural characteristics of the galler species (body length of the adult, development time, overwintering site, oviposition location (i.e. oviposition in the leaf blade, through the leaf midrib, or elsewhere), and phenology). Gall type was obtained from the literature (Kopelke et al. 2017), whereas other gall characteristics were measured in the field. Adult body lengths were measured on collected specimens by author AL, and remaining traits were obtained through expert consultation.

Parasitoid traits included physical traits (ovipositor length and body length), and behavioural traits (whether it was an endoparasitoid or ectoparasitoid, the stage at which it attacked the galler, reproduction (sexual or asexual), and phenology). Most traits were obtained from the literature (Kopelke et al. 2017), whereas physical traits were measured on collected specimens, and remaining traits were again obtained through expert consultation.

Some traits were highly correlated, so we removed them from our analysis. Specifically, we removed tree height, gall toughness, overwintering site, reproduction strategy of the galler, parasitoid/inquiline, koinobiont/idiobiont, and parasitoid attack strategy. For further details, see the Supporting information. After removal of correlated traits, we had eight *Salix* traits, eight galler traits, and five parasitoid traits.

For species and traits for which we were unable to obtain specific values, we imputed the trait value using available information on related species. For species where we had other species from the same genus (or family), we imputed the trait value as the mean of that trait across other species in the genus (or family). To taxa for which we had no references from the same family or genus, we conservatively assigned the mean value of that trait across all species. *Salix* hybrids were given the mean value of their parent species. Such an approach means that trait values for imputed species will have little impact on the model fit (providing no information beyond the mean), and may therefore mask the importance of such traits. Additionally, it is likely to create a tighter relationship between phylogeny and trait values than may otherwise exist. For details on which traits were imputed, and for how many species, as well as for methods used in reconstructing the phylogenies of each trophic level, see the Supporting information.

Statistical approaches

To dissect the relationship between species traits and network structure, we used a range of mutually complementary statistical approaches, each proposed by recent work on trait matching. Below, we provide an overview of the models used, and of their rationale, whereas for full details on how the models were fitted to empirical data, we refer the reader to the Supporting information. For all analysis we used R (www.r-project.org) and, unless explicitly specified, the methods used in the study and the data on which they have been applied have been implemented and included in the R package 'alien' ver. 1.0 (Blanchet et al. 2020, see the Supporting information for more details).

To compare the different models we relied on Tjur's D (Tjur 2009), a coefficient of discrimination specifically designed for presence-absence data. Tjur's D is calculated by independently averaging the estimated probability for presences and absences of interaction and subtracting the averages for the presences with that of the absences. Tjur's D is asymptotically equivalent to the classic coefficient of determination (R^2) as demonstrated by Tjur (2009), making it an interesting statistic because it can be interpreted as a proportion of variance explained. However, Tjur's D is typically much more sensitive to divergences from the data resulting in values that are much smaller than the classic R^2. A Tjur's D value of 1 represents a perfectly predicted network, while a Tjur's D value of 0 is the null expectation where there is no relationship between traits and network structure. For parametric models (GLM, Fourth corner and IMC), we also calculate AIC. This allows us to ascertain whether additional parameters in the model actually improve the model or not. Unfortunately we cannot calculate AIC for the remaining models, as they do not have a fixed number of parameters.

K-nearest neighbour

The K-nearest neighbour (KNN) algorithm answers the question of whether consumer species with similar characteristics (e.g. similar observed interactions, similar traits, or phylogenetically related species) interact with similar resource species. It works by selecting a species' K closest 'neighbours', i.e. those with the most similar characteristics, and predicting that interaction partners shared by those neighbours are most likely shared also by the focal species. As such, it is a non-parametric method (i.e. the relationship between traits and interactions highlighted in Fig. 1i is determined based on neighbours and not assigned a priori), which allows for a more flexible fit.

We calculated neighbours in three ways: 1) shared interactions, 2) shared traits and 3) shared phylogeny. For shared interactions, we select the K species that share the most observed interactions with the focal species, and use their interactions to predict the interaction partners of the focal species. Although this method may seem circular, in that it uses observed interactions to predict interactions, it tells us about the inherent structure in the data (i.e. to what extent there are groups of species that share similar interactions as their 'neighbours'). As such, it gives the upper limit on how much other metrics of dissimilarity could predict. If the network does not contain groups of species with similar interaction partners, then KNN will perform poorly regardless of what metric of dissimilarity is used (Desjardins-Proulx et al. 2017). We also used similarity in traits and phylogenetic distance as metrics for selecting neighbours. If KNN predictions based on traits and/or phylogeny show a fit comparable to that of KNN using interactions, then the interaction structure of the network is well explained by the traits measured and/or species' relatedness.

To implement the K-nearest neighbour approach, we first constructed a dissimilarity matrix quantifying the relationship between each pair of species within the species pool (Desjardins-Proulx et al. 2017). For this, we focused on 1) observed interactions among species, 2) traits, or 3) phylogeny to calculate the dissimilarity among pairs of species. We used Jaccard dissimilarity (Jaccard 1901) with shared interaction data, Euclidean distance with traits, and cophenetic distance with phylogenetic information. The K species least dissimilar to the focal species were then selected as neighbours. In the current case we used three neighbours ($K=3$) for all KNN analyses. Generally, the number of neighbours to consider in a KNN analysis is explored and the value resulting in the optimal result is adopted for further analyses. In the context of the present study, the large number of missing values in the data (i.e. species never observed to co-occur) prevented us from considering any number of neighbours larger than $K=3$. Note that if fewer than three neighbours could be found for a particular species (because of missing values), then all the available neighbours were considered.

For each pair of species, the sum of interactions for the K nearest neighbours of both species divided by $K \times 2$ was used as the prediction. For example, if two of parasitoid i's three nearest neighbours interacted with galler j, and one of galler j's three nearest neighbours interacted with parasitoid i, the interaction probability (and therefore the likelihood) would be $(2+1)/(3 \times 2)=0.5$.

Random forest

Random forest is a very flexible method and can capture more complex relationships between traits and interactions compared to parametric methods. To determine whether this flexibility is needed, we compared the parametric GLM and Fourth corner results to the flexible Random forest method. Random forest divides the data successively based on those particular trait values that give the best separation between categories (i.e. interactions versus co-occurrence, Fig. 1ii), thus creating a decision tree. To protect against overfitting, Random forest repeats this multiple times to create a 'forest' of decision trees and then uses the mean prediction of all trees (Breiman 2001). Random forest is the most flexible of the approaches used by us. Hence, compared to all other approaches here that use real traits, we expect it to give the upper prediction limit. We apply Random forest using both traits and (separately) phylogeny.

Random forest is, per definition, a univariate technique. Thus, as a first step we converted the matrix describing interactions to a long vector, with corresponding trait values used as explanatory variables. To predict interactions, we generated 2000 trees while defining the minimum size of the terminal nodes as 1 and setting the number of trait (or phylogenetic) variables sampled as candidates at each split to

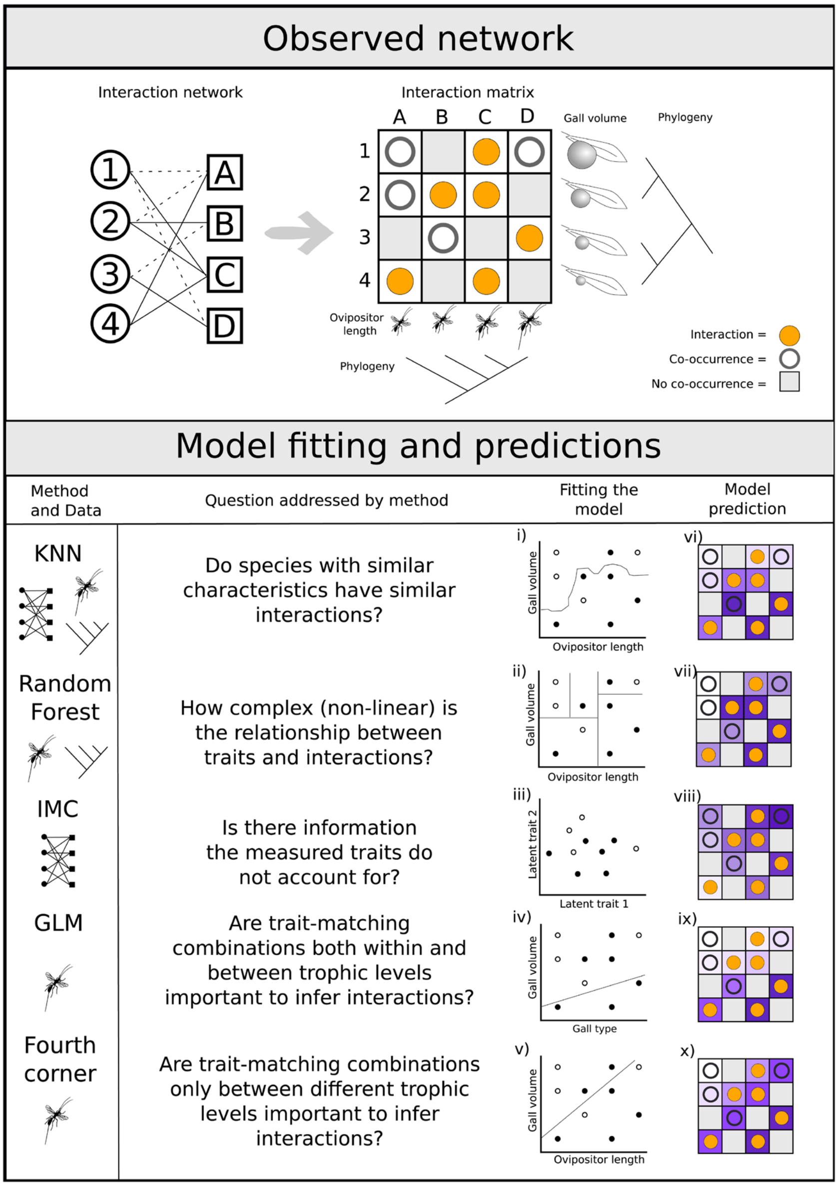

Figure 1. Conceptual overview of our approach, and of the information provided by each method. For each bipartite element in our tri-trophic network, our observations concern a set of consumer species attacking a set of resource species (shown by the 'Interaction network'). We then describe the observed interactions by a matrix of the resource species (numbers) and consumer species (letters), where a yellow filled dot denotes an interaction, an open dot describes a co-occurrence without interaction, and a grey cell describes two species that never co-occurred at the same site. For concreteness, we here illustrate the galler–parasitoid network. Each species is associated with a set of traits, here represented by gall type and ovipositor length (but our data actually include many more traits) and phylogeny. As illustrated under 'Model fitting and predictions', we fit each model to the data (i–v), generating a prediction (vi–x) where darker purple colors specify a higher probability that species may interact. To fit the model, we used different data sources to inform link prediction, as indicated by the icons under 'Method and data'; we used i) network structure (shown by the network icon) for KNN and IMC, ii) trait values (shown by the parasitoid icon) for all methods except IMC and iii) phylogeny (shown by the phylogeny icon) for KNN and Random forest. Note that in the figure we present overly simplified examples to convey a conceptual understanding. As such, we illustrate analysis using only two traits. Similarly, panel (ii) illustrates the fitting of a Random forest, and presents only a single decision tree for visual clarity. Full models are presented in the Statistical approaches.

$\sqrt{p}$ (the default value), where p is the total number of variables. The Random forest analysis were carried out using the alien package, which is effectively a wrapper for the R package `'randomForest'` (Liaw and Wiener 2002). As such, the default values were used for the arguments not discussed in the previous lines. We are aware that for optimal performance, further parameter tuning may be valuable depending on the data the model is applied on. However, we found that even when using the default values, Random forest performed well and could provide valuable insights about the structure of our networks. As the aim of this study is to show that using complementary methods together can help us gain a deeper understanding of our system, we therefore chose to apply the same adjustment across the two different datasets.

For Random forest, we also constructed phylogenetic variables in addition to (but separately from) traits. These phylogenetic variables were derived from the cophenetic matrix. Specifically, we used all axes obtained from a principal coordinate analysis (PCoA) (Gower 1966) as calculated on the cophenetic matrix to evaluate the value of phylogeny in identifying interactions.

Indirect matching centrality

Indirect matching centrality (IMC) is a parametric method that uses the structure of the data to provide the upper limit on how much parametric methods using real traits could be expected to explain. For this purpose, we use it as the baseline to compare generalized linear models and the Fourth corner analysis.

The IMC framework uses latent traits to optimally describe the structure of the food web, as based on the match between predator and prey (latent) traits and the centrality of each species (Rohr et al. 2010, 2016). 'Latent' traits are simply values assigned to each species to capture as much variation in the data as possible (Fig. 1iii). They may, but do not necessarily, correspond to any real traits. Species are arranged along a centrality dimension (describing a species' propensity for making interactions) and a matching dimension (describing whom they are likely to interact with). These matching and centrality dimensions may correspond to one or a combination of real traits. Yet, real traits can maximally explain as much as a latent trait and generally significantly less (Eklöf et al. 2013, Rohr et al. 2016). Specifically,

$$\mathrm{Logit}(P(L_{i,j}=1)) = -\sum_{k=1}^{d}\lambda_k\left(m_{1,i}^{k}-m_{2,j}^{k}\right)^2+\delta_1 c_{1,i}+\delta_2 c_{2,j}+m \qquad (1)$$

where $L_{i,j}$ is the interaction between species i and j. The model is defined to have d matching dimensions – here we used $d=1$ – but only one centrality dimension. $m_{1,i}^{k}$ and $c_{1,i}$ are respectively the kth (out of d) matching and centrality latent traits associated with one set of species (e.g. predator, indicated by the subscript 1) while $m_{2,j}^{k}$ and $c_{2,j}$ are the kth matching and centrality latent traits associated with the other set of species (e.g. prey, subscript 2). λ_k is the kth matching parameter weighting the importance of matching, while δ_1 and δ_2 are parameters weighting the importance of centrality for species set 1 and 2 respectively. With only a single matching dimension as we use here, a higher value of λ than δ corresponds to a more modular network, while the opposite corresponds to a more nested network (Rohr et al. 2010, 2016). Lastly, m is the model intercept. To make it easier to relate the equation above to the description of the method by Rohr et al. (2010), we use the same notation.

Generalized linear models

GLMs are traditionally one of the most common methods for capturing the importance of trait matching. As a parametric model, the functional form between traits and interactions is defined a priori (Eq. 2) and it is therefore less flexible than KNN and Random forest. Note that GLM (and Fourth corner) can be understood as matching-centrality methods using real traits instead of latent traits (for the latter, see IMC, above), or 'direct matching-centrality' methods.

Because the interaction matrix characterized the presences and absences of each interaction, we used a quadratic logistic regression model with interactions. Specifically, the logistic regression model is defined as

$$\begin{aligned}\mathrm{Logit}(P(L_{i,j}=1)) &= \beta_0+\beta_1 T_{1,i}+\beta_2 T_{2,i}+\beta_3 T_{1,j}\\ &+\beta_4 T_{2,j}+\beta_5 T_{1,i}^2+\beta_6 T_{2,i}^2+\beta_7 T_{1,j}^2+\beta_8 T_{2,j}^2\\ &+\beta_9(T_{1,i}\times T_{2,i})+\beta_{10}(T_{1,i}\times T_{1,j})+\beta_{11}(T_{1,i}\times T_{2,j})\\ &+\beta_{12}(T_{2,i}\times T_{1,j})+\beta_{13}(T_{2,i}\times T_{2,j})+\beta_{14}(T_{1,j}\times T_{2,j})\end{aligned} \qquad (2)$$

where $L_{i,j}$ is the interaction between species i and j, $T_{1,i}$ and $T_{2,i}$ are traits associated to the set of species i (e.g. predator), and $T_{1,j}$ and $T_{2,j}$ are traits associated to the set of species j (e.g. prey). The subscripts 1 and 2 simply identify two different traits. All βs are the model's estimated parameters, which were used to calculate the prediction. (Note that only continuous traits were squared, not categorical traits.)

Fourth corner

The Fourth corner approach attempts to use the three 'corners' of 1) the interaction matrix, 2) the matrix of resource species × their traits, and 3) the matrix of consumer species × their traits to predict the 'Fourth corner' of resource species' traits × consumer species' traits. Fourth corner then explains how interactions between resource and consumer traits, i.e. trait matching, affects the interactions recorded in the interaction matrix. In practice, GLM and Fourth corner differ only by the combination of traits used. Specifically, the Fourth corner analysis only includes trait matching combinations between species of different trophic levels, while GLM includes combinations both within and between levels. For example GLM may include the trait combination ovipositor length × oviposition strategy, but Fourth corner would not (because they are both parasitoid traits). Therefore, we compared the two methods to determine whether trait matching within a trophic level is important in addition to matching between levels. Fourth corner can at best perform as well as GLM, and hence the difference between the two tells us the importance of trait matching within a trophic level (Fig. 1v). In applying the Fourth corner approach, our interest was in studying the interaction matrix using the two sets of traits as explanatory variables. As such, the model was defined as follows:

$$
\begin{aligned}
\mathrm{Logit}(P(L_{i,j}=1)) &= \beta_0 + \beta_1 T_{1,i} + \beta_2 T_{2,i} \\
&+ \beta_3 T_{1,j} + \beta_4 T_{2,j} + \beta_5 T_{1,i}^2 + \beta_6 T_{2,i}^2 \\
&+ \beta_7 T_{1,j}^2 + \beta_8 T_{2,j}^2 + \beta_9 (T_{1,i} \times T_{1,j}) \\
&+ \beta_{10} (T_{1,i} \times T_{2,j}) + \beta_{11} (T_{2,i} \times T_{1,j}) + \beta_{12} (T_{2,i} \times T_{2,j})
\end{aligned}
\tag{3}
$$

The notation used in the above equation is the same as for GLM. As for GLM, only continuous traits were squared. However, contrary to GLM, the Fourth corner method accounts for interactions solely between traits of species of different trophic levels. In other words, e.g. the term $\beta_{14}(T_{1,j} \times T_{2,j})$, representing an interaction between trait 1 and trait 2 of species group j, occurs in the GLM model (Eq. 2) but not the Fourth corner model (Eq. 3).

We chose the five methods introduced above to provide ecological perspective on our data, but stress the availability of additional methods not used here. These include Ives and Godfray (2006)'s phylogenetic analysis of trophic interactions, methods relying on co-occurrence (Harris 2016) and other machine learning algorithms (Pichler et al. 2020). Methods based on latent variables show particular promise in accounting for sampling biases and unobserved interactions (Papadogeorgou et al. 2023).

Comparing subsets of traits

When all traits were used concurrently, it was difficult to determine which traits contributed the most to model fit. We therefore compared subsets of four traits at a time. By comparing which traits were frequent in subsets that performed well (i.e. based on Tjur's D), we can see which traits contributed the most to the relationship between traits and network structure. Additionally, this allowed us to compare models using the same traits – a solution not possible when using all traits, because GLM and Fourth corner were not able to use all traits concurrently (Supporting information). We ran each analysis that uses real traits (KNN using traits, Random forest using traits, GLM, and Fourth corner) with all possible combinations of four traits, where at least one trait came from the resource trophic level and one from the consumer trophic level (to allow for matching between consumer and resource traits). We then calculated and compared model performance for each analysis across each trait combination.

Results

In summary, methods using network structure explained approximately twice as much structure as equivalent methods using species traits or phylogeny (KNN with interactions versus KNN with traits or phylogeny, and IMC versus GLM and Fourth corner, Table 1, Fig. 2). However, all KNN models fit very poorly to the *Salix*–galler part of the network (Fig. 3, Supporting Information). The more flexible, non-parametric Random forest method performed the best of all methods using traits for both networks, and performed better with phylogeny than traits. Gall type was the most important trait identified for GLM and Fourth corner (Fig. 4), while Random forest performed best with continuous traits, and no traits stood out strongly for KNN. Overall, this suggests that our traits – especially gall type – explain a substantial portion of network structure, but certainly not all, and that more flexible models are more adept at capturing this relationship. These overall results are outlined in further detail below.

K-nearest neighbour methods

In the galler–parasitoid network, KNN using interactions gave a good fit to the data (Table 1, Fig. 3). KNN using phylogeny or traits explained roughly half as much as KNN with interactions in the galler–parasitoid network (Table 1). Thus, the galler–parasitoid network is structured into groups of gallers/parasitoids interacting with similar parasitoid/gallers and this structure may be partially – but not wholly – due to these species having similar traits or being phylogentically closely related. In the *Salix*–galler network however, KNN was by far the worst-performing analysis, even when using interactions to determine neighbours (Table 1). This suggests that the *Salix*–galler network does not have this structure. Indeed, the *Salix*–galler network is composed primarily of highly specialized galler species, making it difficult to find neighbours.

Table 1. Log-likelihood, Akaike information criterion (AIC) and Tjur's D for the different analyses applied to the *Salix*–galler and galler–parasitoid data set. Note that the GLM and Fourth-corner methods have constraints that force them to use a subset of traits (Supporting information). Also, because in KNN and Random forest analyses the number of parameters cannot be evaluated, AIC cannot be calculated. For the log-likelihood and AIC, values closer to zero are better. For the Tjur's D, higher values are better.

	Salix–galler			Galler–parasitoid		
Model	log-likelihood	AIC	Tjur's D	log-likelihood	AIC	Tjur's D
KNN interactions	−812		−0.07	−742		0.59
KNN phylogeny	−563		−0.01	−1217		0.25
KNN traits	−731		−0.06	−1213		0.30
RF traits	−143		0.42	−500		0.63
RF phylogeny	−89		0.63	−278		0.76
IMC	−135	286	0.53	−854	1724	0.45
GLM traits	−268	830	0.22	−1163	2673	0.27
Fourth-corner traits	−279	782	0.20	−1223	2728	0.22

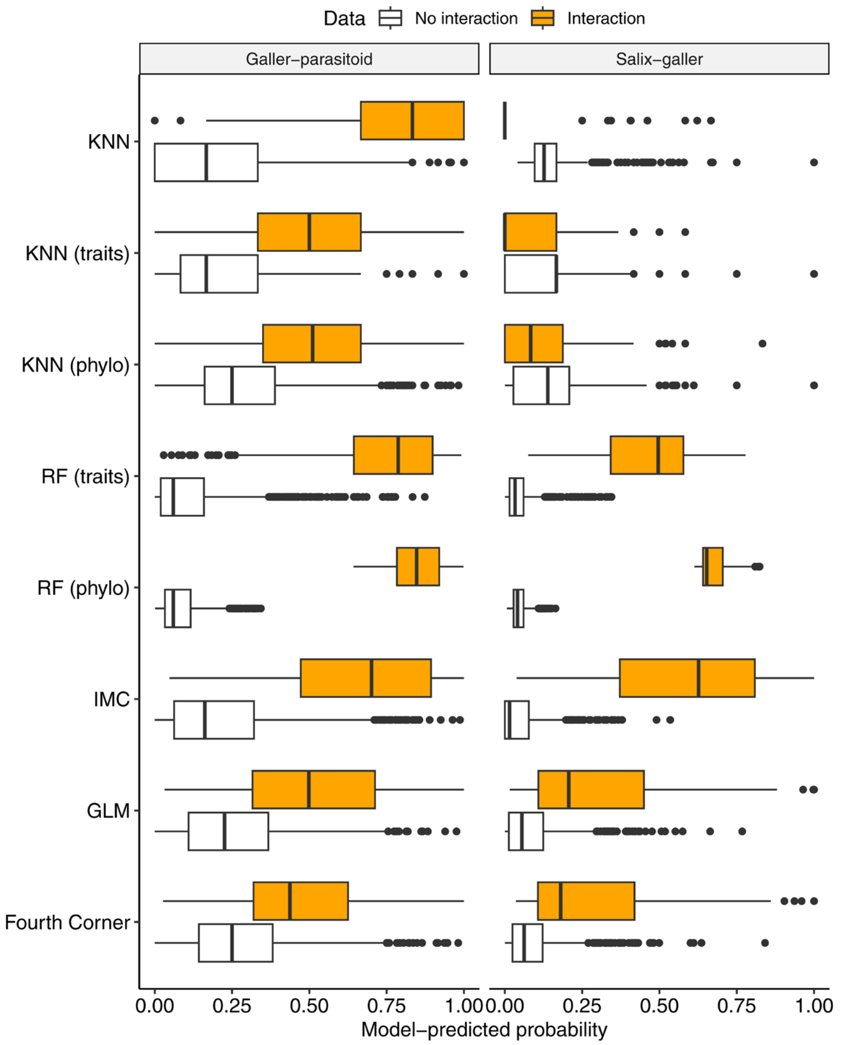

Figure 2. Probability of interaction as predicted by each model, separated by interactions which did exist in the data (yellow) and those that did not (white), for both the galler–parasitoid network (left panel) and the *Salix*–galler network (right panel). Better performing models have yellow boxes closer to one (predicting a high interaction probability for interactions that did exist) and white boxes closer to zero (low predicted interaction probability for interactions that did not exist). Boxes represent the median and first and third quartiles of model-predicted interaction probabilities.

Indirect matching-centrality

We used IMC to determine the upper limit of how much of network structure a parametric method could explain. For the galler–parasitoid network, IMC performed roughly as well as KNN using interactions, while IMC performed substantially better than KNN for the *Salix*–galler network (Table 1). The performance of IMC suggests that the networks can be well characterized using only two dimensions (one for matching and one for centrality). From the fitted parameter values that scale the relative importance of matching (λ in Eq. 1) versus centrality (δ in Eq. 1) dimensions, we can see that the *Salix*–galler network ($\lambda^{SG} = 2.5$, $\delta_G^{SG} = 1.23$, $\delta_S^{SG} = 2.5$) relies on the matching (λ) dimension to a greater degree than the galler–parasitoid network ($\lambda^{GP} = 0.74$, $\delta_G^{GP} = 1.42$, $\delta_P^{GP} = 1.38$). Here, we note that the centrality dimension of gallers for the *Salix*–galler network (δ_G^{SG}) is much weaker than the centrality dimension of *Salix* species. This is due to a high proportion of galler species being strict specialists (i.e. only interacting with 1–2 species) in their interactions with *Salix*, but not with parasitoids.

GLM, Fourth corner and Random forest

GLM, Fourth corner and Random forest all use real traits to predict interactions. By comparing them to IMC, which uses latent traits, we can thus determine how well the traits explicitly sampled really explain network structure. At the same time, we can make a comparison between the parametric IMC, GLM and Fourth corner methods on the one hand, and the more flexible Random forest method on the other hand, to thereby determine whether the relationship between traits and interactions is well-explained by the structure of

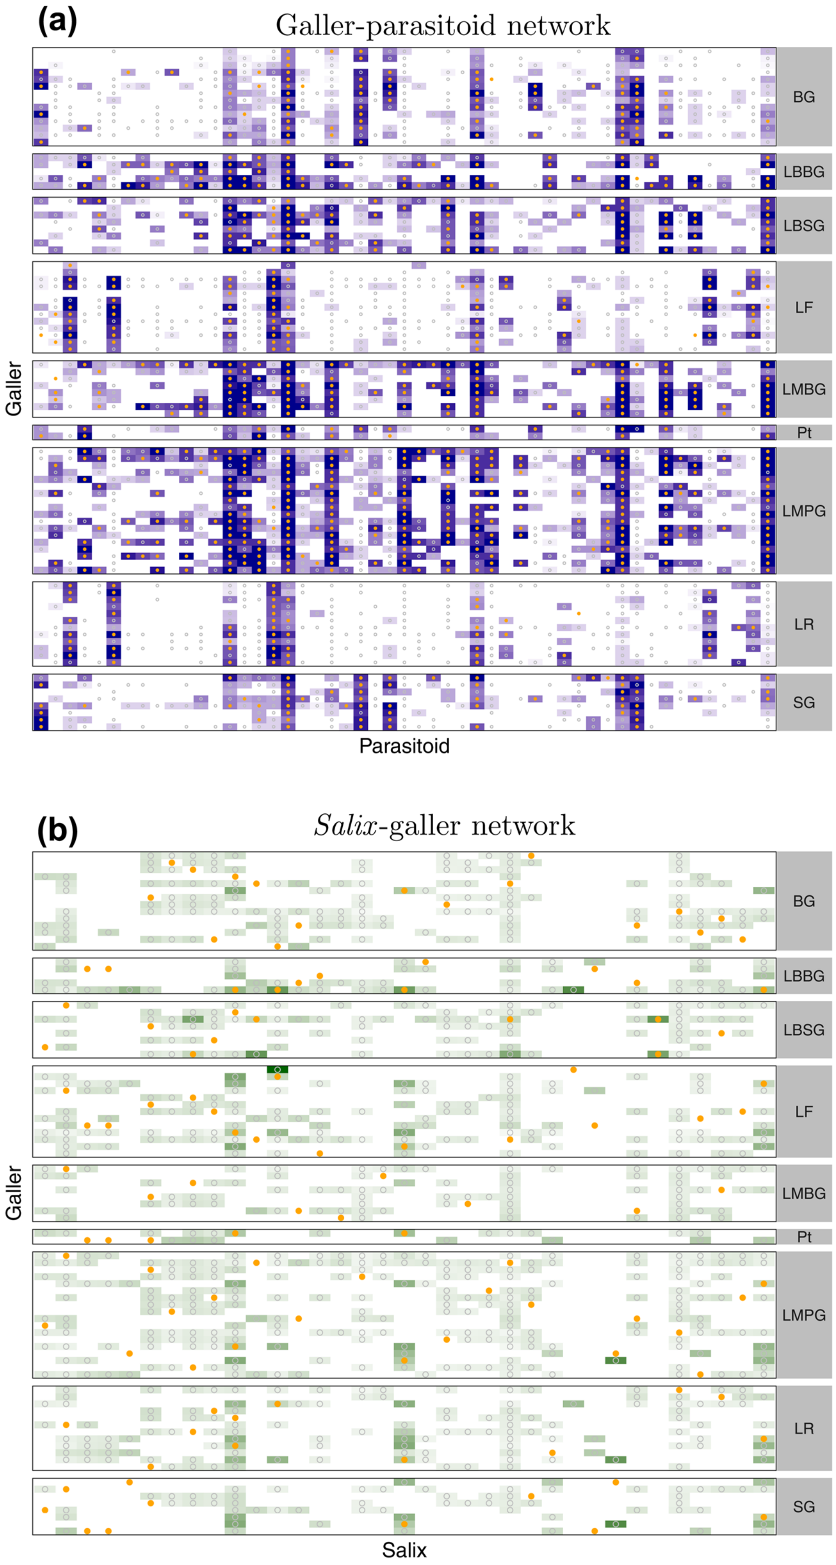

Figure 3. Galler–parasitoid (a) and *Salix*–galler (b) interaction matrices showing model predictions from K-nearest neighbour using shared interactions to calculate dissimilarity. Each column represents a parasitoid species (in (a)) or *Salix* species (in (b)) and each row represents a galler species. Observed interactions are shown by yellow points, while gray points indicate co-occurrence but not interaction. A missing point refers to a lack of cooccurrence. Model predictions are shown in purple or green, with darker shades specifying the prediction of more likely interaction. Gallers are grouped according to gall type: LF = leaf folder, LR = leaf roller, BG = bud galler, SG = stem galler, Pt = petiole galler, LBBG = leaf blade bean gall, LMBG = leaf midrib bean gall, LBSG = leaf blade sausage gall, LMPG = leaf midrib pea gall. See the Supporting information for similar figures of other methods, and for the names of the species in this plot.

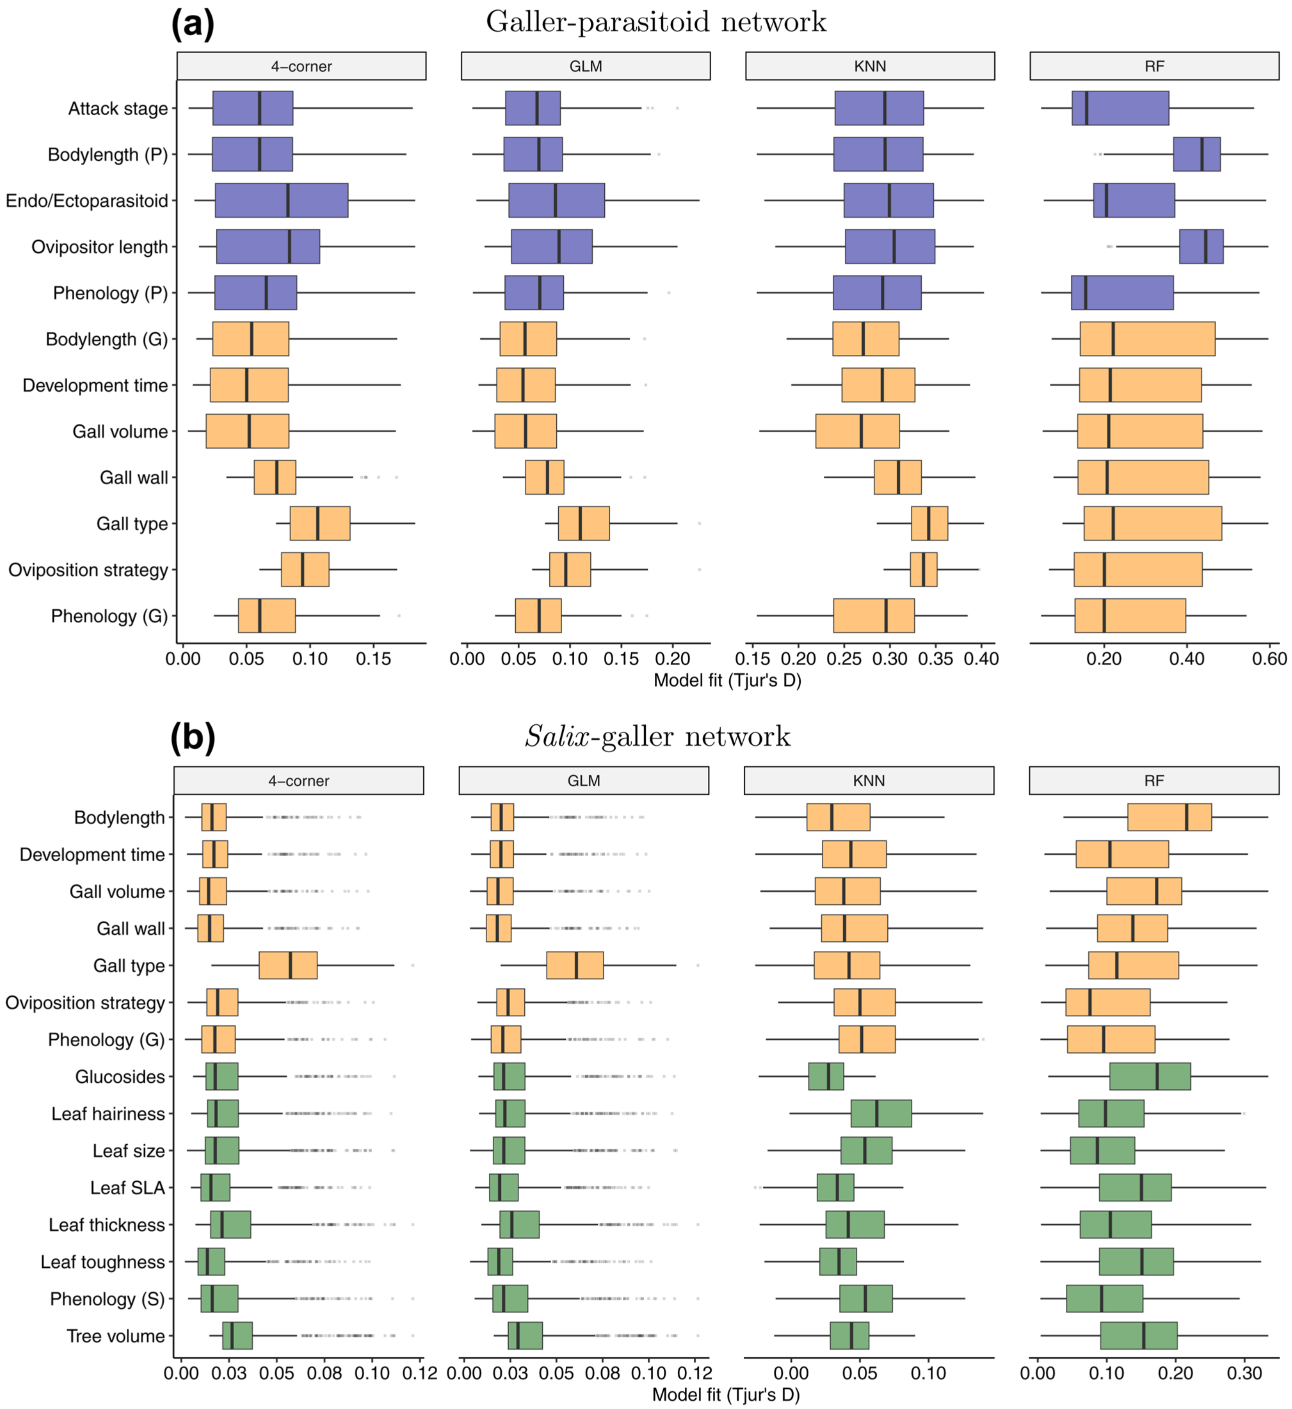

Figure 4. Model fit (Tjur's D values) for trait-based methods using combinations of four traits. Boxes represent the median and first and third quartiles of model fits for traits combinations including the named trait. Traits with consistently higher values of Tjur's D therefore contributed more to model performance. *Salix*, galler and parasitoid traits shown in green, yellow, and purple respectively. Note differences in x-axes scales.

the parametric models or whether a more flexible method is needed.

When considering traits, we found that Random forest performed roughly equally to IMC (Table 1). At the same time, for both networks, GLM performed roughly half as well as IMC, with Fourth corner close behind (Table 1). The strong performance of Random forest suggests that the nature of the relationship between traits and interactions has complexities better captured by the more flexible method, while the performance of GLM and Fourth corner relative to IMC tells us that our measured traits only account for roughly half of the potentially explainable network structure.

The difference in fit could be partly attributed to the fact that GLM and Fourth corner use a smaller subset of traits than Random forest (Supporting information). Despite this, it seems that the more flexible Random forest method better captured the complexities of the relationship between traits and network structure than did GLM or Fourth corner methods, with a caveat: Random forest could only capture these complexities given sufficient flexibility, i.e. when using continuous variables, which can be split at any point, rather than categorical variables, which are limited to the number of categories in the variable. Meanwhile, the difference in performance between GLM and Fourth corner was greater in the galler–parasitoid than *Salix*–galler network, implying that interactions between traits at the same trophic level may play a greater role in the galler–parasitoid network than in the *Salix*–galler network.

Random forest using phylogeny was the best-performing method for both networks (Table 1) – a finding consistent with the greater correlation observed between phylogeny and interactions than between traits and interactions for almost all species groups (Supporting information).

Correlation between interactions, traits and phylogeny

We examined the correlations between the interaction, trait and phylogenetic dissimilarity matrices for species within the same trophic level. To do so, we calculated the Pearson's correlation between pairs of matrices that had been converted to vectors. This is another way of establishing whether species with similar traits (or species that are more closely related) interact with similar species at other trophic levels. In both data sets (*Salix*–gallers and galler–parasitoids), we observed higher correlations between species' interaction similarity, trait similarity and phylogenetic similarity among resource species than among consumer species (Supporting information). The correlations were strongest between phylogenetic similarity and interaction similarity.

Correlations between methods

To examine the consistency in predictions between different models, we calculated the correlation in the probability of interaction assigned to each species pair by the different methods. These correlations were uniformly high for the galler–parasitoid data (Table 1, Supporting information), despite important differences in likelihood. With one exception, the strongest correlations were between the best-performing models; Random forest using phylogeny, and Random forest using traits, KNN using interactions and IMC. The weakest correlations in predictions occurred between GLM and Fourth corner with other methods, although they still showed relatively strong correlation with other methods using traits. The strongest correlation in both networks was between GLM and Fourth corner, although this is unsurprising given closely similar model structures.

The range of correlations proved much more diverse between methods for the *Salix*–galler than the galler–parasitoid data (Table 2, Supporting information). After the strong correlation shown between GLM and Fourth corner, the highest correlations were among the KNN models, which predicted similar structures (but did so poorly). However, KNN methods had very low correlation with all other methods. Predictions for the three best-performing methods for the *Salix*–galler network, two Random forest analyses and IMC, were tightly correlated, although notably lower than in the galler–parasitoid network.

In a nutshell, we found methods using network structure, real traits, and phylogeny to be mostly in agreement about which interactions are most or least likely to occur. This applies in particular to the galler–parasitoid network, where the resource–consumer interactions are not as specialised as in the *Salix*–galler network.

Which traits explain the most?

To establish the importance of individual traits, we identified traits which were frequently associated with improved model fit in combinations of only four traits. To this aim, we compared the performance of individual traits among the methods which explicitly focused on real traits (KNN with traits, GLM, Fourth corner and Random forest), observing a wide range of variation in values of Tjur's D (Fig. 4).

Firstly, we observe that the type of trait, in particular, whether it was categorical or continuous, can strongly affect model performance. Overall, in both networks, Random forest was the best-performing method, however, Random forest also frequently performed poorly and, for certain combinations of traits, worse than GLM (Supporting information). This occurred when using categorical traits with few

Table 2. Correlation between model predictions when using all traits (or phylogeny) for the galler–parasitoid data (upper triangle, in purple) and *Salix*–galler data (lower triangle, in green). All correlations are significant, with $p < 0.001$.

	KNN	KNN (tr)	KNN (ph)	RF (tr)	RF (ph)	IMC	GLM	Fourth corner
KNN	1	0.667	0.591	0.775	0.813	0.752	0.529	0 512
KNN (traits)	0.772	1	0.628	0.701	0.634	0.671	0.596	0.582
KNN (phylo)	0.782	0.719	1	0.624	0.574	0.599	0.518	0.481
RF (traits)	0.070	0.142	0.159	1	0.876	0.709	0.631	0.603
RF (phylo)	0.153	0.185	0.213	0.697	1	0.772	0.547	0.510
IMC	0.130	0.152	0.177	0.494	0.584	1	0.554	0.496
GLM	0.007	0.051	0.051	0.461	0.271	0.231	1	0.879
Fourth corner	0.020	0.064	0.067	0.425	0.257	0.189	0.921	1

categories – such as whether the parasitoid was an endo- or ectoparasitoid – (Supporting information), and can be attributed to the fact that these traits constrain the flexibility of the Random forest approach. Categorical traits can be split at most into the number of categories in the data, while continuous traits can be split in infinite ways. The linear GLM is much less affected by this difference.

Some models were more correlated in their predictions than others, when using the same subsets of four traits. Here, GLM and Fourth corner were almost perfectly correlated in their predictions for both networks, but showed weak correlation (galler–parasitoid network) or no correlation (*Salix*–galler network) with predictions from Random forest (Supporting information). KNN using traits showed strong correlations with GLM and Fourth corner and a weaker correlation with Random forest for the galler–parasitoid network. This results contrasts with the one obtained from the *Salix*–galler network, where predictions from KNN showed essentially no correlation with predictions from GLM and Fourth corner and a surprisingly strong negative correlation with the prediction from Random forest (Supporting information).

At the level of specific traits, gall type was consistently identified as the most important trait by Fourth corner and GLM in both networks and by KNN for the galler–parasitoid network only (Fig. 4). In the galler–parasitoid network for GLM, Fourth corner, and KNN, oviposition strategy (whether the galler oviposits in the leaf blade, leaf vein or elsewhere) came a close second and gall-wall thickness third. For KNN in the *Salix*–galler network, no traits stood out this strongly (Fig. 4) and values of Tjur's D were consistently low, suggesting a systematically poor fit. For Random forest, the traits consistently associated with good fits were traits with continuous values, including body length of both galler and parasitoid, and ovipositor length.

Discussion

Ultimately, all interactions will be structured by traits, since most species will be fundamentally unsuited for interacting with each other, as due to vast differences in size, feeding organs or the like. The great hope of trait-based approaches is then that a few traits will be strongly predictive of who interacts with whom – since if different traits are relevant to each and every interaction, then a trait-based approach will be useless. Based on our analysis of one of the largest systematically compiled data sets of species interactions to date (Kopelke et al. 2017), as here supplemented with new trait data, we find that a limited set of traits is indeed important for structuring interactions. Our key insight from applying a full suite of trait-based methods is that depending on the structure of the network, different methods tell us different components of the story. In brief, we find distinctly different imprints of traits within two coupled, bipartite network elements (*Salix*–galler versus galler–parasitoid interactions). Methods using network structure (KNN using interactions) or latent traits (IMC) rather than real traits gave a good fit to galler–parasitoid interactions and (for IMC only) *Salix*–galler interactions. These results show that interactions are arranged in a way that could potentially be explained by a limited set of traits. Whether or not we have then measured the right traits to explain this structure is less clear. The traits used here explain a major part of detectable structure, while clearly not capturing all variation. In the galler–parasitoid data, methods using empirically-established trait values explained roughly half of the network structure explained by equivalent methods not using traits. Methods using traits performed less well for the *Salix*–galler network (explaining roughly 60% less than equivalent methods not using traits). This suggests either that traits were less important for interactions in this network or, more likely, that we have collected information on the wrong traits. Methods using phylogeny performed very well in the galler–parasitoid network and, for Random forest with phylogeny, in the *Salix*–galler network. Using *Salix*-based networks as a model system, our results thus reveal how the relationship between species traits, network structure, and phylogenetic history can be understood through a combination of complementary analyses. Below, we will go through the role of different traits for different types of interactions, and the insights shed by the different methods.

Similar species are more likely to share interactions (in the galler–parasitoid network)

We first asked whether similar species are more likely to share interactions. Based on the KNN analyses, we find that this is indeed the case in the galler–parasitoid subnetwork, but not the *Salix*–galler subnetwork. The poor performance of KNN in the latter network is due to a combination of highly specific interactions (most gallers only interacted with 1 or 2 *Salix* species) and the fact that many *Salix* and galler species never co-occured. Together, this results in most species having few neighbours and suggests that KNN is a poor method for such networks.

In the galler–parasitoid network, KNN using traits or phylogeny explained approximately half of the network structure explained by KNN using interactions. Of the traits used by us, gall type was the most important trait for explaining this structure, as closely followed by the trait of oviposition strategy. The remaining structure, unexplainable by our traits or phylogeny, may be driven by other, unmeasured traits, which may not be phylogenetically conserved.

Gall type is the most important trait structuring interactions

Gall type was the trait most responsible for the structure detected by KNN, but also for GLM and Fourth corner (Fig. 4). This makes sense if we explore this from an ecological perspective. Since the gall provides a defensive structure, interactions between gallers and parasitoids will depend, in large part, on the parasitoid's ability to penetrate the gall (Stone and Schönrogge 2003). Thus, traits allowing a parasitoid to breach the gall wall are likely important in structuring interactions and, if the parasitoid can access one gall, it can presumably access other, similar galls too. Intriguingly, in the galler–parasitoid, but not the *Salix*–galler network, the trait 'oviposition strategy' was nearly as successful as gall type for explaining network structure. For a galling insect, oviposition strategy can essentially be seen as a looser grouping of gall type, and largely corresponds to how difficult it is to access the gall; leaf rollers and folders oviposit on the leaf blade; leaf blade sausage galls, leaf midrib pea gallers and bean gallers all oviposit through the leaf midrib (Kopelke 1999); and bud and shoot gallers oviposit into the bud and shoot, respectively. From the perspective of a parasitoid, a leaf fold is likely much the same as a leaf roll, a bud gall is likely as difficult to access as a stem gall, and leaf blade sausage galls and leaf midrib pea and bean galls all have a roughly similar structure and gall wall to penetrate (Stone and Schönrogge 2003). As a result, the galler–parasitoid network had a nested structure, where some parasitoids showed many interactions. This structure was well explained by the traits measured as well as by phylogeny, because traits such as gall type are phylogenetically conserved (Nyman et al. 2000, 2007). This structure is also captured by IMC, where both the matching and centrality dimensions show similar importance for the galler–parasitoid network.

While our results imply that there are likely other important traits that we have missed, they are also encouraging. What they suggest is that a limited set of traits can explain a substantial portion of network structure, and that trait matching is indeed an important mechanism structuring networks. From a practical perspective, the success of gall type is particularly hope-inspiring; in terms of using traits to predict interactions, the ideal traits are those that are easily measured (Violle et al. 2007). Gall type is easily categorized and is one of the first pieces of information recorded when describing galler species (Liston et al. 2017). Indeed, it was the only galler trait accompanying the interaction data when the data set behind this paper was originally published (Kopelke et al. 2017). The success of gall type over other traits associated with gall type, such as gall-wall thickness, may be due to the fact that several different traits are important and gall type effectively summarizes them.

The relationship between traits and interactions is best captured by more flexible methods in the galler–parasitoid network…

Multiple traits likely govern a parasitoid's ability to access and parasitize a galler, such as ovipositor length and phenology (Peralta et al. 2020). Indeed, these traits emerged as influential traits in our analysis. These traits, both of the same trophic level, probably interact, as revealed by the better performance of GLM relative to Fourth corner. They likely do so in complex ways, as suggested by the better performance of more flexible methods such as Random forest and KNN. This answers our second question, as to whether common parametric models are sufficient to capture trait matching, or whether more flexible methods are necessary.

… but not the *Salix*–galler network

For the *Salix*–galler network, however, the less flexible IMC model is one of the best performing models, and performs better than Random forest. KNN performs particularly poorly (Table 1). GLM and Fourth corner perform similarly well, suggesting that interactions between traits on the same trophic level are of little importance. An examination of the ecology of this system, however, suggests that the above results are actually due to the importance of 'trait syndromes', where a combination of traits are frequently observed together (Agrawal and Fishbein 2006), driving a tight match between *Salix* and galler species. Many *Salix* species exhibit high intraspecific variation in both phenotypic traits and chemical properties (Skvortsov 1999, Hörandl et al. 2012). Despite this variation, we observed that most galler species are highly specialized with respect to their interactions with *Salix*. Most species only interacted with one, and at most four, *Salix* species. Clearly, there are mechanisms preventing gallers from interacting more widely, and these mechanisms are likely trait-associated. Galls are induced by substances secreted during oviposition (Kopelke 1999, Yamaguchi et al. 2012) and depend on a tight match of galler traits with the *Salix* species, in order for the galler to successfully manipulate the plant into producing a gall. Such traits of the gallers involve e.g. ovipositor morphology (Liston et al. 2017), the chemistry of phytohormones released during oviposition (Yamaguchi et al. 2012) and oviposition behaviour (Roininen et al. 1999). Furthermore, *Salix* species utilize a variety of chemical compounds as protection against herbivory (Volf et al. 2015); such highly toxic or unique secondary metabolites frequently prevent unspecialized insects from feeding on the plants producing them (Agrawal and Fishbein 2006), while specialists evolve mechanisms to overcome relevant defenses (Denno et al. 1990, Treutter 2006, Roslin and Salminen 2008). These trait syndromes and interactions between traits makes it easier to arrange species along a foraging or vulnerability axis, which can then be captured by a single 'matching' latent trait (Rohr et al. 2010, Eklöf et al. 2013), thereby explaining the strong performance of IMC.

Theoretically, trait syndromes should also be captured by traits. A priori, we would then expect a stronger performance of GLM over Fourth corner, and both performing closer to the upper limit given by IMC. However, we had only coarse information on *Salix* chemistry, and sparse information on galler oviposition traits, which together were clearly insufficient to explain the tight match between *Salix* and galler. Phylogeny proved a strong performer, with Random forest using phylogeny providing substantially the best fit of all methods, suggesting that many of these traits are phylogenetically conserved. Indeed, gall type, the most important trait, is phylogenetically conserved, and we found that related *Salix* species tended to interact with similar gallers (Supporting information) and related gallers interacted with similar parasitoids (Supporting information).

Related resource species tended to interact with similar consumers, but not vice versa

The finding that similar *Salix* species tended to interact with similar gallers also formed part of our ecologically perhaps most intriguing finding – that the predictive power of phylogeny proved asymmetric: while related resource species tended to interact with similar consumers, related consumers showed no particular tendency to interact with similar resources. This pattern was supported by several complementary observations, and strikingly repeated across trophic levels. In terms of the *Salix*–galler network, we observe a high correlation between *Salix* species' phylogenetic history and their interaction similarity (Supporting information). This implies that related *Salix* species tend to be exploited by the same galler species. Conversely, no similar correlation occurred between the phylogenetic history of galler species and their interaction dissimilarity when interacting with *Salix* (Supporting information). In terms of the galler–parasitoid network, we again found a correlation between galler species' phylogenetic history and their dissimilarity when interacting with parasitoids, whereas the opposite was not true (Supporting information).

The patterns observed resound with previous observations of related resources sharing similar interactions but not related consumers (Bersier and Kehrli 2008, Naisbit et al. 2012, Eklöf and Stouffer 2016). As a driving force, this type of pattern may arise because consumers compete directly for shared prey. Such processes are likely to be a stronger driver for divergence than the apparent competition of resources sharing predators (Naisbit et al. 2012). Indeed, Rossberg et al. (2006) found that foraging traits need to evolve faster than vulnerability traits to produce food webs like the ones we find in nature, and consumer shifts to new hosts has been proposed as a major mechanism leading to speciation (reviewed by Berlocher and Feder 2002). A consumer that can develop traits allowing it to match the chemical, morphological, or physical traits of a new host can gain access to new enemy- and competitor-free space. If some individuals develop traits better suited to the new host – but this change results in lower suitability to the old host – then these genetic tradeoffs in performance can eventually lead to divergent natural selection (Nyman et al. 2007, Nyman 2010, Leppänen et al. 2014). For gall-inducing sawflies on their *Salix* hosts, there is genetic evidence to support this type of host-associated divergence (Leppänen et al. 2014). Together, these findings suggest that host-associated differentiation (Leppänen et al. 2014) has driven the development of the *Salix*–galler–parasitoid network that we see today.

Conclusions

In conclusion, the results of the five different methods here applied add credence to the hope invested in trait-based approaches. They suggest that a limited set of traits will be informative for understanding who interacts with whom across large networks. They also point to what traits are most important, and what traits we might preferentially measure. Most intriguingly, though, they reveal how species' traits, network structure, and phylogenetic history actually interact and drive each other. Through complementary evidence from the models combined, we thus learnt not only about the current ecology of this ecosystem, but also how it came to be and why the two components of the network are structured as they are. Future applications of our approach will likely yield an even deeper understanding of this and other systems by accounting for not only the presence or absence of interactions, but also their intensity, strength and frequency. Aside from IMC, all other analyses considered in this study can be used with such data.

Acknowledgements – We thank Ged Malsher, Francesca Santaniello and Carol Högfeldt for help in the field. Open access publishing facilitated by University of Canterbury, as part of the Wiley – University of Canterbury agreement via the Council of Australian University Librarians.

Funding – DG was supported by a QCBS program for working groups, DG and GB by an NSERC Discovery Grant, TR and KW by VR grant 2016-04580.

Author contributions

Kate L. Wootton: Conceptualization (equal); Data curation-Lead, Formal analysis (lead); Investigation (lead); Methodology (lead); Project administration (lead); Software (equal); Visualization (lead); Writing – original draft (lead); Writing – review and editing (lead); **F. Guillaume Blanchet:** Conceptualization (supporting); Formal analysis (supporting); Funding acquisition (equal); Methodology (supporting); Software (equal); Writing – original draft (supporting); Writing – review and editing (supporting). **Andrew Liston:** Investigation (supporting); Methodology (supporting); Writing – review and editing (supporting). **Tommi Nyman:** Data curation (supporting); Investigation (supporting); Methodology (supporting); Writing – review and editing (supporting). **Laura Riggi:** Investigation (supporting); Methodology (supporting); Writing – review and editing (supporting). **Jens-Peter Kopelke:** Resources (supporting). **Tomas Roslin:** Conceptualization (equal); Funding acquisition (equal); Supervision (equal); Writing – original draft (supporting); Writing – review and editing (supporting). **Dominique Gravel:** Conceptualization (equal); Formal analysis (supporting); Funding acquisition (equal); Software (supporting); Supervision (equal); Writing – original draft (supporting); Writing – review and editing (supporting).

Transparent peer review

The peer review history for this article is available at https://publons.com/publon/10.1111/ecog.07028.

Data availability statement

Data are available from Github: https://github.com/TheoreticalEcosystemEcology/alien (Blanchet et al. 2021). Novel code is available from Github: https://github.com/kwootton/gallers-paper-code.

Supporting information

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

References

Agrawal, A. A. and Fishbein, M. 2006. Plant defense syndromes. – Ecology 87: S132–S149.

Allesina, S. and Pascual, M. 2009. Food web models: a plea for groups. – Ecol. Lett. 12: 652–662.

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

Bascompte, J., Jordano, P., Melián, C. J. and Olesen, J. M. 2003. The nested assembly of plant–animal mutualistic networks. – Proc. Natl Acad. Sci. USA 100: 9383–9387.

Berlocher, S. H. and Feder, J. L. 2002. Sympatric speciation in phytophagous insects: moving beyond controversy? – Annu. Rev. Entomol. 47: 773–815.

Bersier, L. F. and Kehrli, P. 2008. The signature of phylogenetic constraints on food-web structure. – Ecol. Compl. 5: 132–139.

Blanchet, F. G., Cazelles, K., Gravel, D. and Vissault, S. 2021. alien: comparative analysis of interaction inference methods. – R package ver. 1.0-0, https://github.com/TheoreticalEcosystemEcology/alien.

Boukal, D. S. 2014. Trait- and size-based descriptions of trophic links in freshwater food webs: current status and perspectives. – J. Limnol. 73: 171–185.

Breiman, L. 2001. Random forests. – Mach. Learn. 45: 5–32.

Brousseau, P., 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.

Cohen, J. E. and Newman, C. M. 1985. A stochastic theory of community food webs: I. Models and aggregated data. – Proc. R. Soc. B 224: 421–448.

Dehling, D. M., Töpfer, T., Schaefer, H. M., Jordano, P., Böhning-Gaese, K. and Schleuning, M. 2014. Functional relationships beyond species richness patterns: trait matching in plant–bird mutualisms across scales. – Global Ecol. Biogeogr. 23: 1085–1093.

Denno, R. F., Larsson, S. and Olmstead, K. L. 1990. Role of enemy-free space and plant quality in host-plant selection by willow beetles. – Ecology 71: 124–137.

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

Eklöf, A. and Stouffer, D. B. 2016. The phylogenetic component of food web structure and intervality. – Theor. Ecol. 9: 107–115.

Eklöf, A., Jacob, U., Kopp, J., Bosch, J., Castro-Urgal, R., Chacoff, N. P., Dalsgaard, B., de Sassi, C., Galetti, M., Guimarães, P. R., Lomáscolo, S. B., Martín González, A. M., Pizo, M. A., Rader, R., Rodrigo, A., Tylianakis, J. M., Vázquez, D. P. and Allesina, S. 2013. The dimensionality of ecological networks. – Ecol. Lett. 16: 577–583.

Gibert, J. P. and DeLong, J. P. 2017. Phenotypic variation explains food web structural patterns. – Proc. Natl Acad. Sci. USA 114: 11187–11192.

Gower, J. C. 1966. Some distance properties of latent root and vector methods used in multivariate analysis. – Biometrika 53: 325–338.

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.

Gravel, D., Albouy, C. and Thuiller, W. 2016. The meaning of functional trait composition of food webs for ecosystem functioning. – Phil. Trans. R. Soc. B 371: 20150268.

Gray, C., Figueroa, D. H., Hudson, L. N., Ma, A., Perkins, D. and Woodward, G. 2015. Joining the dots: an automated method for constructing food webs from compendia of published interactions. – Food Webs 5: 11–20.

Harris, D. J. 2016. Inferring species interactions from co-occurrence data with Markov networks. – Ecology 97: 3308–3314.

Hörandl, E., Florineth, F. and Hadacek, F. 2012. Weiden in Österreich und angrenzenden Gebieten, 2nd edn. – Eigenverl. d. Arbeitsbereiches Ingenieurbiologie u. Landschaftsbau, Inst. für Landschaftsplanung u. Ingenieurbiologie, Univ. für Bodenkultur, 164p.

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

Jaccard, P. 1901. Etude de la distribution florale dans une portion des Alpes et du Jura. – Bull. Soc. Vaudoise Sci. Nat. 37: 547–579.

Jonsell, B. and Karlsson, T. 2000. Flora Nordica, vol. 1. – The Bergius Foundation. The Royal Swedish Academy of Sciences.

Julkunen-Tiitto, R. 1989. Phenolic constituents of *Salix*: a chemotaxonomic survey of further Finnish species. – Phytochemistry 28: 2115–2125.

Kopelke, J.-P. 1999. Gallenerzeugende Blattwespen Europas – Taxonomische Grundlagen, Biologie und Ökologie (Tenthedinidae: Nematinae: *Euura, Phyllocolpa, Pontania*). – Senkenbergische Naturforschende Gesellschaft.

Kopelke, J., Amendt, J. and Schönrogge, K. 2003. Patterns of interspecific associations of stem gallers on willows. – Divers. Distrib. 9: 443–453.

Kopelke, J. P., Nyman, T., Cazelles, K., Gravel, D., Vissault, S. and Roslin, T. 2017. Food-web structure of willow-galling sawflies and their natural enemies across Europe. – Ecology 98: 1730.

Laigle, I., Aubin, I., Digel, C., Brose, U., Boulangeat, I. and Gravel, D. 2018. Species traits as drivers of food web structure. – Oikos 127: 316–326.

Leppänen, S. A., Malm, T., Värri, K. and Nyman, T. 2014. A comparative analysis of genetic differentiation across six shared willow host species in leaf- and bud-galling sawflies. – PLoS One 9: e116286.

Liaw, A. and Wiener, M. 2002. Classification and regression by randomForest. – R News 2: 18–22.

Liston, A. D., Heibo, E., Prous, M., Vårdal, H., Nyman, T. and Vikberg, V. 2017. North European gall-inducing *Euura* sawflies (Hymenoptera, Tenthredinidae, Nematinae). – Zootaxa 4302: 1–115.

Lucas, T. C. D. 2020. A translucent box: interpretable machine learning in ecology. – Ecol. Monogr. 90: e01422.

Meier, B., Shao, Y., Julkunen-Tiitto, R., Bettschart, A. and Sticher, O. 1992. A chemotaxonomic survey of phenolic compounds in Swiss willow species. – Proc. R. Soc. Edinb. B 98: 229–232.

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

Naisbit, R. E., Rohr, R. P., Rossberg, A. G., Kehrli, P. and Bersier, L.-F. 2012. Phylogeny versus body size as determinants of food web structure. – Proc. R. Soc. B 279: 3291–3297.

Nyman, T. 2010. To speciate, or not to speciate? Resource heterogeneity, the subjectivity of similarity, and the macroevolutionary consequences of niche-width shifts in plant-feeding insects. – Biol. Rev. 85: 393–411.

Nyman, T., Widmer, A. and Roininen, H. 2000. Evolution of gall morphology and host-plant relationships in willow-feeding sawflies (Hymenoptera: Tenthredinidae). – Evolution 54: 526–533.

Nyman, T., Bokma, F. and Kopelke, J. P. 2007. Reciprocal diversification in a complex plant–herbivore–parasitoid food web. – BMC Biol. 5: 49.

Papadogeorgou, G., Bello, C., Ovaskainen, O. and Dunson, D. B. 2023. Covariate-informed latent interaction models: addressing geographic & taxonomic bias in predicting bird–plant interactions. – J. Am. Stat. Assoc. 118: 2250–2261.

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

Peralta, G., Vázquez, D. P., Chacoff, N. P., Lomáscolo, S. B., Perry, G. L. W. and Tylianakis, J. M. 2020. Trait matching and phenological overlap increase the spatio-temporal stability and functionality of plant–pollinator interactions. – Ecol. Lett. 23: 1107–1116.

Pichler, M., Boreux, V., Klein, A., 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.

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.

Rohr, R. P. and Bascompte, J. 2014. Components of phylogenetic signal in antagonistic and mutualistic networks. – Am. Nat. 184: 556–564.

Rohr, R. P., Scherer, H., Kehrli, P., Mazza, C. and Bersier, L. F. 2010. Modeling food webs: exploring unexplained structure using latent traits. – Am. Nat. 176: 170–177.

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

Roininen, H., Price, P. W., Julkunen-Tiitto, R., Tahvanainen, J. and Ikonen, A. 1999. Oviposition stimulant for a gall-inducing sawfly, *Euura lasiolepis*, on willow is a phenolic glucoside. – J. Chem. Ecol. 25: 943–953.

Roslin, T. and Salminen, J.-P. 2008. Specialization pays off: contrasting effects of two types of tannins on oak specialist and generalist moth species. – Oikos 117: 1560–1568.

Rossberg, A. G. 2008. Part-whole relations between food webs and the validity of local food-web descriptions. – Ecol. Complex. 5: 121–131.

Rossberg, A. G., Matsuda, H., Amemiya, T. and Itoh, K. 2006. Food webs: experts consuming families of experts. – J. Theor. Biol. 241: 552–563.

Schneider, F. D., Scheu, S. and Brose, U. 2012. Body mass constraints on feeding rates determine the consequences of predator loss. – Ecol. Lett. 15: 436–443.

Skvortsov, A. K. 1999. Willows of Russia and adjacent countries. – Taxon. Geogr. Rev. 39: 1–307.

Staniczenko, P. P., Kopp, J. C. and Allesina, S. 2013. The ghost of nestedness in ecological networks. – Nat. Commun. 4: 1391.

Staniczenko, P. P. A., Smith, M. J. and Allesina, S. 2014. Selecting food web models using normalized maximum likelihood. – Methods Ecol. Evol. 5: 551–562.

Stone, G. N. and Schönrogge, K. 2003. The adaptive significance of insect gall morphology. – Trends Ecol. Evol. 18: 512–522.

Tjur, T. 2009. Coefficients of determination in logistic regression models – a new proposal: the coefficient of discrimination. – Am. Stat. 11 63: 366–372.

Treutter, D. 2006. Significance of flavonoids in plant resistance: a review. – Environ. Chem. Lett. 4: 147–157.

Violle, C., Navas, M.-L., Vile, D., Kazakou, E., Fortunel, C., Hummel, I. and Garnier, E. 2007. Let the concept of trait be functional! – Oikos 116: 882–892.

Volf, M., Hrcek, J., Julkunen-Tiitto, R. and Novotny, V. 2015. To each its own: differential response of specialist and generalist herbivores to plant defence in willows. – J. Anim. Ecol. 84: 1123–1132.

Williams, R. J. and Martinez, N. D. 2000. Simple rules yield complex food webs. – Nature 404: 180–183.

Williams, R. J. and Purves, D. W. 2011. The probabilistic niche model reveals substantial variation in the niche structure of empirical food webs. – Ecology 92: 1849–1857.

Williams, R. J., Anandanadesan, A. and Purves, D. 2010. The probabilistic niche model reveals the niche structure and role of body size in a complex food web. – PLoS One 5: e12092.

Wootton, K. L., Curtsdotter, A., Roslin, T., Bommarco, R. and Jonsson, T. 2023. Towards a modular theory of trophic interactions. – Funct. Ecol. 37: 26–43.

Yamaguchi, H., Tanaka, H., Hasegawa, M., Tokuda, M., Asami, T. and Suzuki, Y. 2012. Phytohormones and willow gall induction by a gall-inducing sawfly. – New Phytol. 196: 586–595.