Predicting direct and indirect non-target impacts of biocontrol agents using machine-learning approaches

Hannah J. Kotula^{1*}, Guadalupe Peralta¹, Carol M. Frost², Jacqui H. Todd³, Jason M. Tylianakis^{1,4}

¹ Centre for Integrative Ecology, School of Biological Sciences, University of Canterbury, Christchurch, New Zealand

² Department of Renewable Resources, University of Alberta, Edmonton, Canada

³ The New Zealand Institute for Plant and Food Research Limited, Auckland, New Zealand.

⁴ Bio-protection Research Centre, School of Biological Sciences, University of Canterbury, Christchurch, New Zealand

* Corresponding author

E-mail: hannahkotula@gmail.com (HJK)

Abstract

Biological pest control (i.e. 'biocontrol') agents can have direct and indirect non-target impacts, and predicting these effects (especially indirect impacts) remains a central challenge in biocontrol risk assessment. The analysis of ecological networks offers a promising approach to understanding the community-wide impacts of biocontrol agents (via direct and indirect interactions). Independently, species traits and phylogenies have been shown to successfully predict species interactions and network structure (alleviating the need to collect quantitative interaction data), but whether these approaches can be combined to predict indirect impacts of natural enemies remains untested. Whether predictions of interactions (i.e. direct effects) can be made equally well for generalists vs. specialists and across different habitat types is also untested for consumer-prey interactions, though previous work on mutualist networks suggests that interactions among generalists may be more difficult to predict. Here, we used two machine learning techniques (random forest and k-nearest neighbour; KNN) to test whether we could accurately predict empirically-observed quantitative host-parasitoid networks using trait, abundance, and phylogenetic information. Then, we tested whether the accuracy of machine-learning-predicted interactions depended on the generality of the interacting partners or on the source (habitat type) of the training data. Finally, we used these predicted networks to generate predictions of indirect effects via shared natural enemies (i.e. apparent competition), and tested these predictions against empirically observed indirect effects between hosts. We found that random-forest models predicted host-parasitoid pairwise interactions (which could be used to predict attack of non-target host species) more successfully than KNN, and this predictive ability depended on the generality of the interacting partners, but not the source (habitat type) of data used to train the models. Further, although our machine-learning informed methods could significantly predict indirect effects, the explanatory power for both direct and indirect interactions was reasonably low. Combining machine-learning and network approaches provides a starting point for reducing risk in biocontrol.

Introduction

Growing public concern about the harmful health and environmental effects of pesticides [1, 2], combined with the rapid evolution of pest resistance to chemical control [3-5], mean that biocontrol agents are increasingly being advocated to suppress pests [6-8]. Although biocontrol agents are often a more environmentally friendly means of suppressing pests than synthetic pesticides, they can have nontarget effects on native species [9-11]. Therefore, tools for assessing the risk of proposed biocontrol agents prior to their release are crucial for deciding whether a given agent should be introduced [12, 13]. However, in addition to directly affecting individual nontarget species (i.e. direct effects), biocontrol agents have the potential to affect communities through indirect effects (i.e. the effect of one species on another mediated by a third species) [14]. These indirect effects are more difficult to observe (or predict) because of the inherent difficulty in studying whole communities of interacting species and attributing causation to changes in potentially indirectly affected populations. Although various tools (e.g. host range experiments) exist for assessing the potential direct effects of proposed biocontrol agents, to our knowledge, there are currently no methods available for quantifying and ranking potential nontarget impacts of proposed biocontrol agents via indirect effects prior to their release [13], despite them being likely to be common [15, 16].

In particular, populations of herbivore species may be linked via their shared natural enemies (e.g. parasitoids, which are often used as biocontrol agents), such that an increase in one host herbivore population may drive an increase in the abundance of a shared enemy and result in a decrease in another host's population (i.e., apparent competition [17]). In this way, parasitoid species may mediate indirect effects, such as apparent competition, among host species. Similarly, herbivores used as control agents of weeds may influence native herbivores via shared enemies. Although typically not quantified in biocontrol risk assessment, many studies indicate that apparent competition, mediated via shared natural enemies, may be important in structuring insect communities [17-21].

Studies assessing apparent competition are mostly focused on its effects between single species pairs, whereas less is known about its potential community-wide impacts. However, quantitative food webs, which include information on the abundance of both consumer and resource species and the frequency of their interactions, can indicate the potential for indirect effects and assist in the generation of testable hypotheses [22, 23]. To test such a hypothesis, a seminal study experimentally removed two leaf-miner species from a community and found that other leaf miners, with which they shared parasitoids, experienced lower parasitism rates [24]. Similarly, Tack *et al.* [25] experimentally increased the abundance of three leaf miner species and, using information on their shared parasitoids, found long-term indirect interactions between some species, as predicted, though these effects were positive (apparent mutualism [17]) rather than negative (apparent competition). These two studies indicate that indirect effects could, in theory, be predicted for all host species in an assemblage. Congruent with this expectation, a recent community-level empirical study found that apparent competition was important for structuring host-parasitoid networks in native and plantation forests, and that these two habitats functioned dynamically (and somewhat predictably) as a single unit [21]. Although the potential negative impacts of apparent competition on native communities have been considered in a biocontrol context [26], it remains difficult to predict these impacts prior to the release of a biocontrol agent.

The analysis of ecological networks is an important tool in pure and applied community ecology [27-29], and the above examples suggest that it offers a promising approach to understanding the community-wide impacts (via direct and indirect interactions) of biocontrol agents [11, 30]. This network framework has been widely used in ecology, including to quantify and predict indirect effects, and less widely in a biocontrol context to successfully evaluate the impacts of biocontrol agents on non-target species through post-release studies [31-34]. Yet, few studies use it to predict biocontrol outcomes prior to the release of an agent (but see [35]), despite the clear benefits of such an

approach [16, 22, 36]. One notable exception was a study by López-Núñez *et al.* [35], which used networks to predict the potential indirect effects of a biocontrol agent. Because networks were only collected prior to the release of the control agent, López-Núñez *et al.* [35] were unable to validate these predictions. Additionally, their predictions of indirect effects were informed by known interaction partners of the control agent (data that may not be available for a given proposed control agent).

One of the limitations to the use of (particularly quantitative) networks in biological control is the cost or effort required to resolve community-scale interactions. A potential solution would be to use frameworks that predict, rather than sample, species interaction networks using existing data on species traits. Species traits play an important role in structuring species interactions and ecological networks [37-40]. For instance, ecological and life history traits determine species distributions, as well as allowing morphological and physiological trait matching between interacting partners, and thereby influence the occurrence and frequency of interactions [38]. When traits cannot be easily measured or the traits responsible for interactions are unknown, phylogenies, which depict the shared evolutionary history of species, can serve as a surrogate for measured and unmeasured traits, and thus help to explain the pattern of interactions within a community [41-44].

Accordingly, many studies have shown that traits and phylogenies can successfully predict direct interactions among species [37, 38, 45, 46]. More recently, machine-learning techniques, such as random forest and k-nearest neighbour (KNN), have been added to the list of tools used to predict species' interactions based on their traits and/or phylogenies. These algorithms are trained using known interactions, which allow the machine-learning technique to learn which trait and/or phylogenetic combinations are likely to correspond to an interaction occurrence, and these can then be used to predict interactions among new species [47, 48]. In addition to traits and phylogenies (i.e., information that captures niche processes), neutral processes, which describe interaction patterns

based on the probability of species randomly encountering each other, play an important role in shaping ecological networks [49-51]. This neutral component of interactions can be incorporated into predictive models (including machine-learning approaches) by adding data on species' abundances [52].

Although the above literature suggests that traits, phylogenies and abundances can be useful predictors of species interactions, less is known about the circumstances under which predictions of interactions based on these data are likely to be better (or worse). Emerging research suggests that species characteristics (e.g., trophic generality or biogeographic status; native vs. exotic) may influence the predictive ability of trait- (including phylogenetically-) or abundance- based models [53, 54].

Additionally, research has indicated that predicting network structure at a new habitat can be difficult [51], potentially because species composition and the interactions that are realised, which in turn can influence species' realised generality or abundance, can vary with habitat type [55]. In the context of biological control, it would be necessary to predict interactions for agents in their introduced location (as an exotic species potentially interacting with a non-target native community) and across a diversity of both natural and crop habitats. Thus, understanding the robustness of predictions to habitat or biogeographic differences and trophic generality would be crucial.

Combining machine-learning and network approaches, if successful, could allow non-target impacts of proposed parasitoid biocontrol agents to be predicted given a list of host species in the recipient community, along with known host-parasitoid interactions and species data (traits and phylogenies) to inform a machine-learning model. Known interactions from either a biocontrol agent's native range or the community to which an agent is to be introduced could be used. Alternatively, if predictions of interactions can be made equally well within vs. across habitat types, network data from a different habitat could be used to inform a predictive model.

In this study, we examine whether network approaches can be used to predict effects (both direct and indirect) of species in parasitoid-host networks, with a view to using these approaches in biological control. However, rather than focusing on biocontrol species only, we focus on developing a proof of concept using an entire host-parasitoid community (only some of which are pests and biocontrol agents), which comprises a large-scale dataset, sampled in two different habitat types (plantation and native forest), from Frost *et al.* [21]. Specifically, our objectives were to 1) compare two machine-learning techniques (random forest and KNN) for the prediction of (quantitative) host-parasitoid network interactions (i.e. direct effects), using traits, abundances, and phylogenies to inform the models (which could complement host-range studies currently used to assess direct effects); 2) test whether the ability of these machine-learning techniques to predict species interactions depended on a) species' trophic generality or b) whether the data used to train the predictive model came from the same vs. different habitat type as that for which the predictions were being made; and 3) determine whether these machine-learning-predicted interactions, when combined into networks, can be used to predict indirect effects (specifically apparent competition), which could then be applied to non-target effects of biocontrol agents.

Methods

Study system

We collected quantitative host-parasitoid interaction data from 16 training sites (8 from native forest and 8 from plantation) and 32 testing, i.e. validation, sites (16 from native forest and 16 from plantation) in the Nelson/Marlborough region of New Zealand (further details of sampling methods are provided in Frost *et al.* [21] and Peralta *et al.* [56]). Each training site consisted of 2 transects; an edge (10 m towards the forest interior from the edge between the two forest types, and running parallel to the edge) and interior (400-500 m from the forest edge). Each test site consisted of 1 edge transect (from its respective forest type). The test sites did not include interior transects; however, we chose to include both edge and interior transects in the training sites to use all available data. At each

transect (50 m x 2 m), all foliage-dwelling Lepidoptera larvae ('hosts') were collected from plants and reared in the laboratory until they developed into adults or parasitoids emerged (Hymenoptera, Diptera and Nematoda; though we excluded Nematoda parasitoids from all our analyses as the traits we measured, see 'Traits and phylogenies' below, were not suited to these species). Parasitoids and hosts were then identified to species level where possible using available taxonomic information [57-60] and expert assistance (J. Dugdale, J. Berry, and R. Schnitzler), and otherwise to morphospecies. For parasitoids, morphospecies were validated as species using molecular methods. The training sites were sampled seven times over two summer seasons (season 1: December 2009, January, February 2010; season 2: October, November 2010, January, February 2011). To determine whether predictions of apparent competition were equally effective for large experimental and small natural changes in host abundances, host abundance was experimentally reduced in half of the plantation test sites by aerial spraying with a caterpillar-specific biopesticide in the middle of sampling. Test sites were sampled twice before this host reduction treatment (time t , October, November 2010) and twice after (time $t+1$, January, February 2011). Because Frost *et al.* [21] found that the predictive ability of apparent competition did not differ across the host-spray treatment, we did not include this variable in our analyses.

Traits and phylogenies

We collected traits of host and parasitoid species that have been identified as being useful for predicting interactions: body size, trophic generality (measured as normalised number of food resource species, termed 'normalised degree' in network studies; ND), and biogeographic status (native vs. exotic). Note that ND was measured for parasitoids based on their interactions with hosts and for hosts based on their interactions with their plant hosts, rather than parasitoids. We chose these traits because body size is correlated with many important characteristics [61] and is often the best predictor of species (including host-parasitoid) interactions [45], while species generalism tends to be conserved across native and alien ranges [62], and recent evidence suggests that niche processes for species interactions (and therefore predictions based on traits and phylogenies) are

weaker for generalists than specialists, and for exotic species than native species [53]. Additionally, abundance (a neutral process) and phenology play a role in determining interactions [63-65]; we therefore calculated these characteristics for both host and parasitoid species. Hereafter, we refer to body size, trophic generality, biogeographic status, abundance and phenology collectively as 'traits', because these characteristics that can drive niche and neutral processes were all treated similarly in analyses. Because phylogenies can also help to explain interaction patterns [44, 46] by capturing variation in difficult to measure traits, we constructed host and parasitoid phylogenies to include along with the traits in the prediction of interactions (see 'Traits and phylogenies' in Appendix S1).

Predicting direct effects

We compared two machine-learning techniques (random forest and KNN) for the prediction of quantitative host-parasitoid interactions (i.e. direct effects). Training site data were used to inform the models and test site data (at time t) were used to test (i.e. validate) the models.

To predict interactions at new locations (or for new species such as biocontrol agents), the random-forest model learns (from the training data) which host and/or parasitoid trait (or phylogenetic) combinations are likely to correspond to an interaction occurrence. The algorithm then uses these rules, along with the traits (or phylogenetic position) of new species, to predict the probability of interaction occurrence between new combinations of host and parasitoid species pairs. These interaction probabilities could be used in risk assessment if it can be demonstrated that they correspond with higher observed interaction frequencies (which we test below). Random forest therefore assumes that the strength of trait (or phylogenetic) matching between species determines, in part, the likelihood of an interaction between them. In contrast to random forest, which incorporates matching of parasitoid traits (or phylogenetic position) to those of hosts, KNN uses the host-use preferences of parasitoid species (from the training data) with similar traits (or phylogenetic position) to recommend (i.e., predict) interaction partners and interaction frequencies for new parasitoid species. This is analogous to Netflix recommending films to users based on the preferences

of similar users. Thus, in KNN, parasitoid traits are used to determine similarity among parasitoid species, rather than to match with host traits as in random forest (see 'Machine-learning techniques' in Appendix S1).

Model variants with different training data

When making predictions about the impacts of potential biocontrol agents, it is possible that interaction and/or trait data will not exist for the specific recipient habitat. Therefore, to determine whether the ability of a machine-learning technique (random forest or KNN) to predict interactions differed when training data came from the same vs. different habitat than the test habitat, we built three models ('native', 'plantation', 'combined') for each technique, using data (i.e. pooled 'meta-networks') from the training sites of each habitat category. We created one host-parasitoid training meta-network for each of the three models, by pooling data across training sites (of a single forest type, native or plantation, for the 'native' and 'plantation' models, respectively, and all training sites, i.e., both forest types, for the 'combined' model) and sampling dates (see 'Training meta-networks' in Appendix S1). Because the random forest was predicting the probability of an interaction occurring, whereas KNN predicts a weight of the interaction (i.e. frequency), we trained the random-forest models on the binary (occurrence vs. not of each interaction) meta-networks and trained the KNN models on the weighted meta-networks (though for both approaches a quantitative interaction network was predicted).

For KNN, recommendations (i.e., predictions) are based on host use by parasitoid species in the training data, so only host species present in the training data will be recommended to parasitoid species in the test data. In contrast, host species only found in the test data, hereafter called 'new' host species, will never be recommended to a parasitoid - these host species are analogous to new films that no one has seen or rated yet. To overcome this problem, we included 'new' host species into the training meta-networks (only for KNN models because random forest does not suffer from this

problem) by recommending parasitoid species from the training data to each 'new' host species, using KNN. In this case, we recommended parasitoid species to hosts by using host traits and phylogenetic position to determine their nearest neighbours (i.e. the host species most similar to them) (see 'KNN – adding 'new' host species' in Appendix S1).

We tested the predictions of each model on all test sites where parasitoids emerged from any collected host at the first-time step, t (i.e., before the host spray treatment: 14 sites from native forest and 15 sites from plantation). Test sites at which no parasitoids emerged at time t were excluded from all analyses ($N=3$). For each test site, we predicted an interaction frequency (for KNN; see 'KNN – predicting interactions at test sites' in Appendix S1) or a probability of interaction occurrence (for random forest) between each host-parasitoid species pair present at the site. If a parasitoid species was found in both the training and testing data, it was not used to inform predictions about its own interaction partners in the KNN algorithm (i.e. it was not used as a neighbour of itself); this decision was made to be conservative, because interaction data in the new location would likely not be available for a proposed control agent.

We performed a grid search to find the optimal hyperparameters (including which traits to include) for the random forest and KNN models (see 'Tuning hyperparameters' in Appendix S1). Random forest automatically weights traits by their importance (i.e. less informative traits contribute less to predictions). To determine which traits were most important for predicting host-parasitoid interactions using random-forest models, we used the 'feature importances' function from Python's Scikit-learn v 0.23.1 package [66]. In contrast to random forest, in KNN all traits contribute equally to measures of species' similarity. For this reason, unimportant traits must first be removed, as retaining them in the model can decrease predictive ability. In our case of predicting host-parasitoid interactions using KNN models, we found that body size, phenology, parasitoid ND (normalised degree; a measure of generality), and abundance were useful parasitoid traits, whereas phylogeny was removed during our

parameter search process (see ‘Best hyperparameters for KNN’ in Appendix S1). Additionally, we found that all host traits except abundance (i.e., body size, phenology, phylogeny, host ND and biogeographic status) were useful (and thus retained in the models). For the random forest models, parasitoid traits in decreasing order of importance were abundance, phenology, ND, phylogeny, body size, while host traits in decreasing order of importance were phenology, phylogeny, body size, abundance, ND, and biogeographic status (for both hosts and parasitoids this is the average importance across all three models: ‘native’, ‘plantation’, ‘combined’). All random forest and KNN analyses were performed in Python (version 3.8.3). We used the ‘RandomForestClassifier’ function from the Scikit-learn package [66] to perform the random-forest analyses, and used base Python for KNN.

Predicting indirect effects

For each machine-learning (random forest or KNN) approach, we used the quantitative interaction network predicted at each test site (at time t) by the model trained on both native forest and plantation data (i.e. the ‘combined’ random forest or KNN model, respectively) as a basis for predicting indirect interactions, specifically apparent competition. We used these ‘combined’ models because, in a biocontrol context, available data may not be specific to a target habitat, but rather aggregated from different sites and habitats.

Predicting indirect effects involved three steps. First, we used the ‘combined’ training meta-network (see above) to calculate a regional quantitative measure of shared parasitoids, also known as the potential for apparent competition, d_{ij} , using Equation 6 from Müller *et al.* [23]; see ‘Measure of shared parasitism’ in Appendix S1. This d_{ij} value measures, for every pair of host species in a community, the proportion of parasitoids attacking host species i that recruited from host species j [23]. We used the PAC function from the bipartite v 2.15 R package [67] to calculate d_{ij} . Second, for each machine-learning approach (random forest and KNN), we used this measure of shared parasitism (d_{ij}), along with random forest- (or KNN)-predicted initial attack rates (α in equation 1 in Appendix S1, see also

'Calculating α from predicted networks' in Appendix S1) by each parasitoid on each host from time t at the test sites combined with known changes in host abundances (the change in abundance of each host species from sample t to $t+1$, in the site for which predictions are being made) to calculate the expected parasitism rate (E) at time $t+1$ of host species i for each test site. Parasitism rate is the number of parasitism events divided by the number of hosts sampled. We used Equation 3 from Frost *et al.* [21] to calculate E:

$$E_{i(t+1)} = \sum_{j=1}^H \left(\frac{d_{ij} \sum_{l=1}^P \alpha_{ilt}}{n_{jt}} n_{j(t+1)} \right) \left(\frac{1}{n_{i(t+1)}} \right) \quad (1)$$

where n is host abundance, t is the first-time step of data (before experimental host reduction in half of the sites, with $t+1$ being after), α is the link strength (that is, number of attacks), i and j are a focal host species pair, H is the total number of host species, k is a parasitoid species, and l is all parasitoid species, from 1 to P (the total number of parasitoid species).

Finally, we tested whether this expected parasitism rate significantly predicted the observed parasitism rate of each host species i at time $t+1$ test sites (see equation 2 in Appendix S1 for how observed parasitism was calculated).

The logic of this approach is that a bipartite network of interactions, where links represent attack rates between hosts and parasitoids (the nodes), can be projected into a unipartite network with only hosts as nodes and links representing the proportional sharing of parasitoids (d_{ij}). Changes in abundance of hosts in this unipartite network should induce proportionate (and thus predictable) changes in the abundance of their parasitoids and cause a proportionate change in attack rates on any other hosts that share those parasitoid species (i.e. an indirect effect). Frost *et al.* [21] showed that these expected rates (E) significantly predicted observed parasitism rates; this method for predicting indirect effects has thus been experimentally validated using this dataset. Our aim here was to test whether the

machine-learning-predicted networks (trained on data from different sites) could be used instead of observed networks to predict these indirect effects.

The number of host species for which we were able to calculate an expected parasitism rate at each site was much lower than the number of species collected at each site. This difference occurred because we could only calculate an expected parasitism rate for host species within a site that were collected at both time steps and had a non-zero predicted attack rate (α) at the first-time step (t , see equation 1 above). However, for both machine-learning approaches, we used the potential for apparent competition (via the sharing of parasitoids) with every other host species that was collected in the site, to calculate the expected parasitism rates for these species. Therefore, predictions of expected parasitism rates for these species were based on data from the whole network. The numbers of species within sites for which it was possible to calculate an expected parasitism rate were 91 and 84 for the random forest and KNN approaches, respectively.

Statistical analyses

Can random forest and KNN predict interaction frequencies and does this predictive ability differ for generalists vs. specialists and whether the model was trained on data from the same vs. different habitat?

To test the possibility that random-forest-predicted probability and KNN-predicted-frequency can predict observed interaction frequencies, and to compare the predictive ability of each random forest with that of its corresponding KNN model ('native', 'plantation', or 'combined'), we fitted two Poisson generalised linear mixed effect models using observed interaction frequency (from the test sites at the first time step) as the response variable. We included random-forest-predicted probability as a fixed effect in the first model and KNN-predicted interaction frequency in the second model. Including just these predictors in the models allowed us to compare the two machine-learning methods directly, and we then added further predictor variables as a separate step to test our other hypotheses (see below).

This process was repeated two additional times for a total of six models ('native', 'plantation' and 'combined' for both random-forest and KNN predicted probability and frequency, respectively). For each model, we scaled and centred the fixed effect to help with model convergence. In all full models, we included parasitoid identity as a random factor (such that predicted and observed frequencies for a given parasitoid would be 'paired'), and a random slope for random-forest-predicted probability (or KNN-predicted-frequency), to allow this slope to vary across parasitoid species. We also included site as a random factor, to account for the structural non-independence of interactions across sites, though site explained little variance because it covaried strongly with parasitoid composition. We used parasitoid identity because estimating the random slope required multiple interactions per parasitoid, whereas using host-parasitoid pair (see 'Model set 1' in Appendix S1) would have resulted in a model in which the number of observations was less than the number of random effects.

Fitting the model six times (three times for each machine learning method) increases the probability of Type I error. Although sequential Bonferroni correction is often used to account for performing multiple statistical tests, it has been argued to be too conservative [68, 69]. We therefore used a Bernoulli process to calculate the probability of each fixed effect being significant by chance alone, given the number of tests (six) performed [68]. We use this Bernoulli process (for fixed effects) for all subsequent analyses that involve multiple statistical tests.

The prediction of interaction frequencies (which is needed, in turn, to predict indirect effects) represents a stricter test than the prediction of interaction occurrence, as it requires not only the correct prediction of interactions, but also of their strength. However, if this information is not needed, the machine-learning-predicted probabilities and frequencies can be converted into binary values to predict interaction occurrence, so we also explored such predictions of interaction presence/absence (see 'Tests on the prediction of interaction occurrence' in Appendix S1).

For all analyses, we selected the best-fitting model by first selecting the optimal random-effect structure [70]. We ran the full model (described above), as well as models with all combinations of random effects, and selected the model with the lowest AIC. In one case (the KNN plantation model), the model with the lowest AIC value had a singular fit, so in this case we selected the model with the next lowest AIC value that wasn't singular. Then, using this optimal random-effect structure for each model, we selected the best-fitting fixed effects structure by using the dredge function (from the MuMIn v 1.43.17 package [71]), which calculates AIC for the full model as well as all possible simpler models, thereby avoiding issues associated with stepwise selection of variables.

To compare each random forest with that of its corresponding KNN model, we calculated AIC for each best-fitting model. We primarily use this model selection to identify variables that significantly contribute explanatory power, and caution against the mixing of model selection and frequentist statistics [72]. However, for readers accustomed to P values, these are presented in the results.

To determine whether the ability to predict interaction frequency (using machine-learning models) varied with species generality (measured as normalised degree, ND) or according to whether the model was trained on data from the same vs. different habitat, we fitted the same set of models as described above, but with additional predictor variables. Specifically, in addition to random-forest-predicted probability or KNN-predicted interaction frequency, we also included parasitoid ND and host ND as interacting fixed effects (i.e. this model included all possible interactions among these three fixed effects), along with forest type (of the test site) as a fixed-effect interacting with predicted probability or predicted frequency. Because we did not have a specific hypothesis for the forest type x ND interactions, we did not include these interactions in the models. We followed the same model selection process as described above.

Can we predict indirect effects using machine-learning (random forest and KNN) informed methods?

To determine whether expected parasitism rate can predict abundance-driven changes in the observed parasitism rate (at time $t+1$ test sites), and therefore whether indirect effects can be predicted, we fitted a binomial mixed effect model including observed parasitism rate (a proportion) as the response variable. We included expected parasitism rate as a fixed effect and site as a random effect (to control for the non-independence of interactions at a given site). We did not include forest type (native forest vs. plantation) of the test sites as a fixed effect in this model, because Frost *et al.* [21] found that the ability to predict observed parasitism rates did not differ across this variable in this study system. We scaled and centred the expected parasitism rate fixed effect to help with model convergence.

We fitted a separate version of this model to the predictions of expected parasitism rate (see equation 1 above) from each machine-learning approach. For each version of the model, we followed the same model selection process as described above.

Do predictions of indirect effects differ depending on whether observed vs. predicted network data are used to calculate the potential for apparent competition (d_{ij})?

As an additional supplementary analysis, we tested whether predictions of observed parasitism rates differed depending on the network used to calculate the potential for apparent competition (d_{ij}). The rationale for this analysis was that the observed training data are only a sample of the true network, and thus may miss interactions that are important for driving indirect effects. Thus, it is possible that predicted networks better capture these unsampled interactions. Conversely, it is also possible that errors in the predicted networks could lead to poorer estimation of d_{ij} and resulting predictions of indirect effects. For this analysis, we re-ran the above analyses using d_{ij} values calculated from the machine-learning predicted networks, rather than the observed training networks. Overall, this approach did not improve the prediction of indirect effects, but for completeness we present it in the

supplementary materials (methods and results in Appendix S1: see 'Tests using d_{ij} calculated from predicted network data' in Appendix S1).

Do predictions of indirect effects based on machine-learning-predicted networks differ from those based on observed network data?

To compare the random forest and KNN approaches with each other, and with a purely data-driven approach used by Frost *et al.* [21], we re-ran the above analyses (but not the additional analyses using machine-learning predicted d_{ij} values) using only the subset of host species for which it was possible to calculate an expected parasitism rate with all three methods (random forest, KNN, and data-based as used by Frost *et al.* [21], $N=20$). To calculate expected parasitism rate using the data-based approach used by Frost *et al.* [21], we calculated the initial attack rates (α) from the observed time t networks at the test sites and calculated d_{ij} from the 'combined' training meta-network. This is the same approach used by Frost *et al.* [21], except that we did not use habitat specific d_{ij} values and used both edge and interior sites from the training data (rather than just edge sites). We made these changes to the data used because, in a biological control context, site-specific data are less likely to be available, and also to increase the amount of data used to inform predictions in the machine-learning approaches. Additionally, we did not include the seven host species for which we could not determine their biogeographic status (see 'Biogeographic status' in Appendix S1). As before, we performed model selection for each of the three approaches (random forest, KNN and data-based) to find the best-fitting model following the same process described above, except that we used AICc (which is more suited to small sample sizes than AIC), to select the best-fitting model. To compare the three approaches, we calculated the AICc of the best-fitting model for each approach.

All statistical analyses were performed using R v 4.0.1 [73]. We used the lme4 v 1.1.23 package [74] to fit mixed effects models. We used the r.squaredGLMM function (with the delta method) from the MuMIn v 1.43.17 package [71] to calculate marginal and conditional values of R^2 for the mixed effects

models and the rsquared function from the piecewiseSEM v 2.1.0 package [75] to calculate R^2 for the generalised linear models. For all models we checked for overdispersion, and where present we included an observation-level random factor [76], and in all cases this resulted in a model that was not over-dispersed (i.e., the ratio of the sum of squared Pearson residuals to residual degrees of freedom was <2 and non-significant when tested with a Chi-squared test [77]). We visually inspected all models for outliers and, where present, we calculated the leverage value of all datapoints and removed those with high influence (i.e. outliers [78]). Outliers were present in the models predicting indirect effects, though in all cases the results with the outliers removed were qualitatively the same as the results with the outliers included. Note that we did not remove outliers from the models comparing the three methods (random forest, KNN and data based as used by Frost *et al.* [21]) for predicting indirect effects on the common dataset ($N=20$) to ensure that the number of datapoints was kept constant among models to allow model comparison, though in all cases excluding outliers did not qualitatively change the results.

Results

Can random forest and KNN predict interaction frequencies and does this vary with species' generality and whether the model was trained on data from the same vs. different habitat?

We tested whether machine-learning techniques (random forest and KNN) can predict host-parasitoid interaction frequencies. For the less complex set of models (including only random-forest-predicted probability or KNN-predicted frequency), we found that random forests were more successful than KNN at predicting interaction frequencies. Random-forest-predicted probability was significantly positively related to observed interaction frequency in all three models ('native', 'plantation', 'combined'), and although it captured a reasonable amount of variance in the 'combined' model (marginal $R^2 = 0.168$), it only captured a small amount of variance in the 'native' and 'plantation' models (marginal $R^2 = 0.058$ and 0.097 , respectively), which may not be surprising given the myriad

determinants of interaction frequencies in the field (Fig 1A; Tables 1 and 3 in Appendix S1). In contrast, KNN-predicted interaction frequency was only retained in the 'plantation' best-fitting model, and although it was significant in this model, it explained little variance (marginal $R^2 = 0.017$) (Fig 1B; Tables 2 and 3 in Appendix S1). There was a low probability of the predicted-probability or -frequency fixed effect being significant in 4/6 models by chance ($p < 0.001$, calculated using the Bernoulli process), indicating that the results are unlikely to be due entirely to type I error. For 2/3 models ('native' and 'plantation'), the random-forest model predicted observed interaction frequency better (i.e. lower AIC) than the respective KNN model (Table 3 in Appendix S1). Additionally, for all three models, the random-forest model captured a greater amount of variation in observed frequency (higher conditional R^2 values) than the respective KNN models. All best-fitting models included parasitoid identity as a random factor, and a random slope for random-forest-predicted probability (or KNN-predicted frequency) (except for both the random forest and KNN 'plantation' models which only included parasitoid identity as a random factor).

Fig 1. Predictions of observed interaction frequency.

(A) Random-forest-predicted probability significantly predicted observed interaction frequency for all three random-forest models ('combined', 'native', 'plantation'), whereas (B) KNN-predicted frequency only predicted observed interaction frequency for 1/3 KNN models ('plantation' model) (dashed lines represent non-significant relationships). Deviance residuals of the best-fitting model excluding the predictor on the x-axis are plotted, to better reflect the partial effect of the predictor. Each point represents a host species within a site at the first-time step (t) test sites.

For the second set of models, testing whether the ability to predict interaction frequency differed for generalists vs. specialists and within vs. across habitats, we found that the predictive ability of random forest and KNN depended on the generality (ND) of the interacting partners (Fig 2; Tables 4 and 5 in Appendix S1), but not on the source of training data. All best-fitting models included random-forest-

predicted probability (or KNN-predicted frequency), parasitoid ND and host ND as fixed effects. There was a significant positive relationship between random-forest-predicted probability and observed interaction frequency in all three models (i.e. higher predicted probabilities corresponded to higher observed frequencies), qualitatively the same as for the less complex models (see above). Predicted probability did not interact with any fixed effects in 2/3 random-forest models ('native' and 'combined'), and could therefore be multiplied by a constant (the slope) to predict raw interaction frequencies for applied purposes. However, in the 'plantation' random-forest model (Table 4 in Appendix S1) the slope of the relationship between predicted probability and observed frequency (i.e. the predictive ability of random forest) depended on the ND of the interacting partners. In particular, there was a significant positive three-way interaction between predicted probability, host ND and parasitoid ND in this model, such that when ND of both host and parasitoid were high, the slope of the relationship between predicted probability and observed interaction frequency was more positive than expected based on additivity of the main effects alone. However, the presence of a negative two-way interaction between parasitoid ND and host ND, along with negative host ND and parasitoid ND main effects, meant that interactions involving a single generalist partner would tend to occur at lower frequency and have a weaker relationship between predicted and observed frequency, but these effects would be weakened when both partners were generalists.

Fig 2. Overall, machine-learning predictions of observed interaction frequency varied with the generality (measured as normalised degree; ND) of the interacting partners.

There was a significant interaction between predicted-frequency (or -probability) and generality of the interacting partners (host and parasitoid) for 4/6 models. Two are presented here to demonstrate the main patterns: (A) KNN 'combined' model (which had a significant positive predicted frequency x host ND x parasitoid ND interaction, as did the 'native' and 'plantation' KNN models and the 'plantation' random-forest model), (B) random-forest 'native' model, which had a significant host ND x parasitoid ND interaction (as did the 'combined' random-forest model). Deviance residuals of the best-fitting

model excluding the host and parasitoid ND variables are plotted as colour contours, to indicate whether observed interaction frequencies are higher (red) or lower (blue) than expected based on machine-learning predictions. Each point represents a host-parasitoid species pair present within a site at the first-time step (t) test sites.

Congruent with that of random forest, KNN-predicted frequency was significantly related to observed interaction frequency in all three models (for the less complex models KNN-predicted frequency was only significant in the 'plantation' model; see above). Additionally, similar to the 'plantation' random-forest model, there was a significant positive three-way interaction between KNN-predicted frequency, host ND and parasitoid ND in all three models, along with several negative two-way interactions. These higher order interactions suggest that the relationship between predicted frequency and observed frequency varies with the ND of the interacting partners, and that failure to control for this interaction made KNN-predicted frequency a poor predictor of observed interaction frequency when used alone in a model.

For all random forest and KNN models, the relationship between interaction probability (or predicted frequency) and observed interaction frequency did not differ between the two forest types (i.e. the forest type fixed effect was removed during model selection), implying that predictions of interaction frequency (based on random-forest-predicted probabilities or KNN-predicted frequencies) are equally effective in both forest types and do not depend on the source of training data.

The random effects (parasitoid identity and a random slope for predicted probability to allow this to vary with parasitoid identity) were retained in the best-fitting 'combined' and 'native' random-forest models, while only parasitoid identity was retained as a random effect in the best-fitting 'plantation' random-forest model. The best-fitting KNN models all included only an observation-level random factor to deal with overdispersion. The fixed and random effects captured a reasonable amount of

variance in all three random-forest and KNN models (Table 6 in Appendix S1). The fixed effects (random-forest-predicted probability/KNN-predicted frequency, parasitoid ND and host ND) explained more variation in interaction frequency in 2/3 random-forest models than their respective KNN models (i.e. higher marginal R^2 values), though for two models ('native' and 'plantation') the KNN models captured a greater overall amount of variation (i.e. higher conditional R^2 values).

The above results contrast with the analyses for predicting interaction occurrence (see 'Tests on the prediction of interaction occurrence' in Appendix S1). All three KNN models ('native', 'plantation', 'combined') predicted interaction frequencies better than binary interaction occurrences (models predicting frequencies had higher marginal and conditional R^2 values than their respective, i.e. 'native', 'plantation', 'combined', model predicting interaction occurrences). In contrast, 2/3 random-forest models ('native' and 'combined') predicted interaction occurrences better than interaction frequencies (higher marginal and conditional R^2 values). Overall, this suggests that raw predicted probabilities (or frequencies) are generally more useful (at least in KNN) for predicting interactions than assigning presence vs. absence of an interaction occurrence based on an arbitrary probability/frequency threshold.

Can we predict indirect effects using machine-learning (random forest and KNN) informed methods?

For both random forest and KNN approaches, there was a significant positive relationship between expected parasitism rate and observed parasitism rate (random forest: z value=2.42, $p = 0.0152$, $df=85$; KNN: z value=2.83, $p = 0.0045$, $df=77$) (Fig 3; Table 7 in Appendix S1). The probability of this relationship being significant in both models by chance was very low ($p < 0.001$, calculated using the Bernoulli process). Furthermore, expected parasitism rate captured a reasonable proportion of variance for both models (random forest: marginal $R^2 = 0.122$, KNN: marginal $R^2 = 0.096$) (Table 8 in Appendix S1). The site random effect was retained in both models and also explained a reasonable amount of variance (random forest: conditional $R^2 = 0.383$, KNN: conditional $R^2 = 0.329$).

Figure 3. Predictions of observed parasitism rate.

(A) Random-forest- and (B) KNN- expected parasitism rate both significantly predicted observed parasitism rate. Residuals of the best-fitting model excluding the predictor on the x axis are plotted. Each point represents a host species within a site in the second-time step ($t+1$) test sites. Three and four outliers, respectively, were removed from (A) and (B), though the results with the outliers included were qualitatively the same in both cases.

Do predictions of indirect effects based on machine-learning-predicted networks differ from those based on observed network data?

The three methods (random forest, KNN and data-based, as used by Frost *et al.* [21]) did not differ in their ability to predict observed parasitism rates (difference in AICc < 1) (Table 9 in Appendix S1). In fact, with the reduced dataset (N=20; data that could be used for all three analyses), expected parasitism rates did not even significantly predict observed parasitism rates for the random forest and KNN approaches, contrasting with the results with the full dataset (see above), likely due to reduced statistical power (Table 10 in Appendix S1). Similarly, expected parasitism rate was not retained in the best-fitting model for the data-based approach. The site random effect was retained in all best-fitting models and captured a reasonable amount of variance (conditional $R^2 = 0.344$ for all models).

Discussion

The increasing advocacy of biocontrol agents to suppress pest species, combined with their potential to have non-target impacts [10, 11, 79], necessitates the study of methods for assessing the risk of new organisms prior to their release as biocontrol agents [12, 13]. Here, we examined whether network approaches can be used to predict direct and indirect effects of species (e.g., proposed parasitoid biocontrol agents) by comparing the ability of two machine-learning techniques (random forest and KNN) to predict any interaction in the host-parasitoid community. We found that random

forest was more successful than KNN for predicting quantitative host-parasitoid pairwise interactions (which could be used to predict attack of non-target host species), using information on a limited set of easily-obtainable traits. In addition, we found that this predictive ability varied with the generality (ND) of the interacting partners, but did not depend on the source of training data (i.e., native forest, plantation, or both habitat types combined). Further, we found that machine-learning informed methods could significantly predict indirect effects, though with relatively low explanatory power.

Random forest predicts host-parasitoid interactions better than KNN

The better predictive ability of random forest, compared with KNN (Fig 1), may be explained by the different ‘philosophies’ of the two methods. While random forest incorporates trait and phylogenetic matching, which are important for determining, and therefore predicting, interactions [37-39, 44, 46], KNN assumes that similar parasitoid species interact with similar host species. However, parasitoid species with similar traits may be dissimilar in their host use, for at least two reasons. First, adaptation to different host plants can lead to insect speciation [80-83], and this may ‘cascade’ to higher trophic levels, leading to host-associated differentiation in parasitoids [84-87]. Similarly, host-shifts (including the formation of ‘host races’) can contribute to parasitoid speciation [88, 89], and this parasitoid speciation, by either adaptation or host shifts, may result in closely related parasitoid species that attack different sets of host species. In this way, parasitoid species with the most similar trait values (or that are most closely related) may be dissimilar in their host use, resulting in poor predictions of host use by KNN models. Second, it is possible that parasitoid species deemed as similar, based on the selection of traits we considered (see ‘Traits and phylogenies’ in Methods), differed in an important trait that we did not include (for example, counter-defences against immune responses of hosts [90]). However, this may be unlikely because we included parasitoid phylogenies, in addition to traits, which capture variation in a variety of measured and unmeasured traits.

Species’ generality impacts predictions of direct effects

The relative strength of the niche and neutral processes that generate plant-frugivore interactions has been shown to change with species generalism [53], and we similarly found for host-parasitoid interactions that the ability of machine-learning methods to predict observed interaction frequency depended on the generality (ND) of the interacting partners (Fig 2). Our random-forest models incorporated trait and phylogenetic matching (two indicators of niche processes), and the dependence of the 'plantation' random forest's predictive ability on species generality is congruent with the idea that the strength of niche processes varies with generality [53]. However, because we included host and parasitoid abundance (as well as phenology) as traits in the random-forest models, we cannot attribute changes in predictive ability to changes in niche processes alone – neutral (abundance-driven) forces could also have operated, and abundance can covary with generality. Thus, we limit our conclusions to the predictive ability of random forests and KNN changing with species generality, rather than inferring a change in the relative strength of niche versus neutral processes.

Predicting networks within vs. across habitat types

Overall, predictions of host-parasitoid interaction frequencies (and occurrences; see 'Tests on the prediction of interaction occurrence' in Appendix S1) using machine-learning techniques were equally effective in both habitat types (plantation and native forest). Furthermore, predictions of interaction frequencies (and occurrences, overall) did not depend on the source of the training data, which contrasts with recent research showing that predicting networks across habitats (i.e. using data from one habitat to predict interactions at a different habitat) is difficult [51]. This difficulty is, in part, because interaction frequencies can differ markedly among habitat types (e.g., [91, 92]), arising from both changes in random encounters (due to changes in abundance) and interaction preferences based on niche processes [51]. Congruent with this, Staniczenko *et al.* [51] found that models that captured systematic changes in interaction preferences between different habitat types had the best predictive ability, while models based only on interaction data performed poorly. Interaction preferences are likely to differ more with increasing habitat dissimilarity (e.g., [91]). Network data from a more similar habitat (such as the two forest types we compared here) will therefore likely predict

interactions better than the forest versus agricultural habitats compared by Staniczenko *et al.* [51], as differences in interaction preferences will likely be smaller. Additionally, similar habitats are more likely to contain species with similar values for traits, which would likely increase the predictive ability of trait-based models. Therefore, the similarity in physical structure (canopy height, etc) and high host and parasitoid species overlap between plantation and native forest habitat may explain why we found no difference in predictive ability of interaction frequencies between the two habitat types. Nevertheless, at face value, the ability to predict equally across habitats and with different training data are positive attributes for biocontrol risk assessment, and suggest that a common dataset (e.g., national databases of species interactions) could be used for risk assessment of interaction frequencies (a pre-requisite for predicting density-mediated indirect effects) in a range of habitats.

Predicting indirect effects

We found that, although machine-learning informed models significantly predicted indirect effects (i.e., for both random forest and KNN approaches, expected parasitism rate was a significant predictor of observed parasitism rate) (Fig 3), the explanatory power of both these approaches was reasonably low (approximately 10% of variance). Surprisingly, the prediction of indirect effects was better (higher marginal and conditional R^2 values) than the prediction of direct effects for KNN (Tables 3 and 8 in Appendix S1), though the highest overall marginal R^2 value was observed for the model of direct effects with host and parasitoid generality (ND) included (Table 6 in Appendix S1). Together, these results may suggest that the direct effects that were predicted well irrespective of generality were those that were most important for driving indirect effects. Previous work on pollination networks suggests that interactions with a strong signal of trait matching are more spatiotemporally constant and functionally important [93], so it is feasible that the interactions best predicted by machine-learning methods based on traits are those that are least variable and contribute most to indirect effects. These methods are therefore a long way from being able to make accurate predictions of the full strength of parasitoid-mediated indirect effects. However, they represent a promising starting point, and could potentially be used to rank hosts along a continuum from low to high risk of indirect

effects. These machine-learning approaches, which use predicted networks rather than observed networks (at time t test sites) to calculate initial attack rates, α , (see 'Predicting indirect effects' in Methods), have the potential to capture interactions that are not recorded in the observed networks due to inadequate sampling.

Interestingly, the machine-learning approaches and the data-based approach used by Frost *et al.* [21] did not differ in their ability to predict indirect effects for the common dataset (N=20) of host species, for which it was possible to calculate an expected parasitism rate with all three methods; this lack of difference was possibly due to reduced statistical power. Moreover, we found that the data-based approach did not significantly predict observed parasitism rate, contrasting the results of Frost *et al.* [21]. This is likely because we used slightly different data to Frost *et al.* [21]. In particular, Frost *et al.* [21] used only edge training sites and explicitly included host habitat (i.e., host species found in different habitats were treated as separate species) when calculating d_{ij} (a measure of shared parasitism), whereas we used both edge and interior training sites and did not explicitly consider host habitat. Species composition and interactions differ between edge and interior sites in this system [94, 95]; thus the inclusion of interior sites may have caused the different results. Additionally, we removed seven host species (for which we could not determine their host biogeographic status) from all our analyses, and we did not include nematodes because their traits were not comparable to those of insect parasitoids. Collectively, these differences likely explain the incongruence with the findings of Frost *et al.* [21].

Although we found that machine-learning-informed models could significantly predict indirect effects (discussed above), part of the field validation data arose from natural changes in abundance, rather than manipulative additions or removals of individuals. It is therefore possible that an unmeasured factor is correlated with both parasitoid sharing and changes in host abundance or attack rates. For example, habitat preferences of parasitoid species may co-vary with parasitoid sharing, such that the

sharing of parasitoids (i.e., apparent competition) is not the cause of the model's ability to predict observed parasitism rates, but rather a correlate. Additionally, we predicted expected parasitism rate, which captures changes in attack rates, but not subsequent changes in host population abundances; a prerequisite for demonstrating apparent competition between host species mediated by shared parasitoids. However, previous work on the same study system showed that accurate predictions of changes in host abundances followed from the ability to predict changes in attack rates, and were, in fact, better predicted than changes in attack rates [21].

Further caveats

Ecological networks are difficult to sample and the difficulty of distinguishing true non-interactions from unobserved interactions, which may simply result from inadequate sampling, is widely recognised [64, 96-98]. Although the 'meta-networks' used to train our machine learning models (see 'Training meta-networks' in Appendix S1) comprised the pooled data across training sites and sampling dates, to maximise the resolution of potential host-parasitoid links, we acknowledge that these meta-networks may not capture all realised interactions.

A further caveat to using these approaches in biocontrol is that predicting interactions involving exotic species (e.g., control agents) may be more difficult than predicting interactions among native species. The lack of co-evolutionary history between exotic and native species in the community suggests that trait or phylogenetic matching will be weaker for these interactions. In fact, Peralta *et al.* [53] found that niche processes were weaker, and Coux *et al.* [54] found that neutral processes were stronger, for plant-frugivore interactions involving exotic species than those among native species, suggesting that methods incorporating trait or phylogenetic matching will perform poorly for interactions involving exotic species compared to interactions between native species. For this reason, an approach that blends information on traits and abundances (as we used here) may help to alleviate the problems of predicting more neutrally-driven interactions. In addition, exotic species tend to interact with numerous species [50, 99, 100] and, as discussed above, interactions involving generalists may be

more difficult to predict based on trait matching than those involving specialists [53]. However, this is less of a concern as biocontrol agents are typically chosen because they are specialists on the target pest.

Conclusion

Previous studies show that biocontrol agents can have negative non-target impacts, and stress the importance of considering these effects in biocontrol risk assessment [79]. Our findings show that machine-learning methods offer a promising approach to begin to predict direct effects, and to a similar extent, indirect effects, using information on species traits, abundances and phylogenies. In addition, the ability to predict species interactions, and more specifically non-target impacts of biocontrol agents, was hypothesised to vary with species characteristics and habitat type [53]. Congruent with the first of these assertions, we found that predictions of interaction frequencies (i.e. direct effects) varied with the generality of the interaction partners (host and parasitoid). However, predictive ability did not depend on the habitat type of the training data. This supports the use of a common dataset of species interactions and traits for risk assessment.

Acknowledgements

We thank Anna Eklöf and Alex James for helpful comments on an earlier draft and Barbara Barratt, Rogini Runghen and Pieter Pelser for helpful discussions.

References

1. Pimentel D, Acquay H, Biltonen M, Rice P, Silva M, Nelson J, et al. Environmental and economic costs of pesticide use. *Bioscience*. 1992;42(10):750-60.
2. Tilman D, Cassman KG, Matson PA, Naylor R, Polasky S. Agricultural sustainability and intensive production practices. *Nature*. 2002;418:671-7.
3. Palumbi SR. Humans as the World's Greatest Evolutionary Force. *Science*. 2001;293(5536):1786-90.
4. Brattsten LB, Holyoke CW, Leeper JR, Raffa KF. Insecticide resistance - challenge to pest-management and basic research. *Science*. 1986;231(4743):1255-60.
5. Denholm I, Rowland MW. Tactics for managing pesticide resistance in arthropods: theory and practice. *Annu Rev Ecol Syst*. 1992;37:91-112.
6. van Driesche RG, Bellows TS. *Biological control*. New York: Chapman & Hall; 1996.
7. Pimentel D. *Techniques for reducing pesticide use: economic and environmental benefits*: John Wiley and Sons; 1997.
8. van Lenteren JC, Bolckmans K, Köhl J, Ravensberg WJ, Urbaneja A. Biological control using invertebrates and microorganisms: plenty of new opportunities. *BioControl*. 2018;63:39-59.

9. Louda SM, Kendall D, Connor J, Simberloff D. Ecological effects of an insect introduced for the biological control of weeds. *Science*. 1997;277(5329):1088-90.
10. Lynch LD, Thomas MB. Nontarget effects in the biological control of insects with insects, nematodes and microbial agents: the evidence. *Biocontrol News and Information*. 2000;21(4):117-30.
11. Memmott J. Food webs as a tool for studying nontarget effects in biological control. In: Follett PA, Duan JJ, editors. *Nontarget Effects of Biological Control*. 1 ed. New York: Kluwer Academic Publishers; 2000.
12. van Lenteren JC, Bale J, Bigler F, Hokkanen HM, Loomans AJ. Assessing risks of releasing exotic biological control agents of arthropod pests. *Annual Review Entomology*. 2006;51:609-34.
13. Barratt B, Howarth F, Withers T, Kean J, Ridley G. Progress in risk assessment for classical biological control. *Biological control*. 2010;52:245-54.
14. Roy HE, Lawson Handley L-J. Networking: a community approach to invaders and their parasites. *Functional Ecology*. 2012;26:1238-48.
15. Simberloff D. Risks of biological control for conservation purposes. *BioControl*. 2012;57(2):263-76.
16. Tylianakis JM, Binzer A. Effects of global environmental changes on parasitoid-host food webs and biological control. *Biological Control*. 2014;75:77-86.
17. Holt RD. Predation, apparent competition, and the structure of prey communities. *Theoretical Population Biology*. 1977;12(2):197-229.
18. Holt RD, Lawton JH. The Ecological Consequences of Shared Natural Enemies. *Annu Rev Ecol Syst*. 1994;25:495-520.
19. Morris RJ, Lewis OT, Godfray HCJ. Apparent competition and insect community structure: towards a spatial perspective. *Annales Zoologici Fennici*. 2005;42(4):449-62.
20. van Veen FF, Morris RJ, Godfray HCJ. Apparent competition, quantitative food webs, and the structure of phytophagous insect communities. *Annu Rev Entomol*. 2006;51:187-208.
21. Frost CM, Peralta G, Rand TA, Didham RK, Varsani A, Tylianakis JM. Apparent competition drives community-wide parasitism rates and changes in host abundance across ecosystem boundaries. *Nature Communications*. 2016;7:12644.
22. Willis AJ, Memmott J. The potential for indirect effects between a weed, one of its biocontrol agents and native herbivores: a food web approach. *Biological Control*. 2005;35:299-306.
23. Müller CB, Adriaanse ICT, Belshaw R, Godfray HCJ. The structure of an aphid-parasitoid community. *J Anim Ecol*. 1999;68(2):346-70.
24. Morris RJ, Lewis OT, Godfray HCJ. Experimental evidence for apparent competition in a tropical forest food web. *Nature*. 2004;428(6980):310-3.
25. Tack AJ, Gripenberg S, Roslin T. Can we predict indirect interactions from quantitative food webs?—an experimental approach. *J Anim Ecol*. 2011;80(1):108-18.
26. van Veen FJF, Memmott J, Godfray HCJ. Indirect Effects, Apparent Competition and Biological Control. 3. Netherlands, Dordrecht: Springer; 2006. p. 145-69.
27. Ings TC, Montoya JM, Bascompte J, Blüthgen N, Brown L, Dormann CF, et al. Ecological networks - beyond food webs. *J Anim Ecol*. 2009;78(1):253-69.
28. Memmott J. Food webs: a ladder for picking strawberries or a practical tool for practical problems? *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2009;364(1524):1693-9.
29. Bascompte J. Structure and dynamics of ecological networks. *Science*. 2010;329(5993):765-6.
30. Ollivier M, Lesieur V, Raghu S, Martin J-F. Characterizing ecological interaction networks to support risk assessment in classical biological control of weeds. *Current Opinion in Insect Science*. 2020;38:40-7.
31. Henneman ML, Memmott J. Infiltration of a Hawaiian Community by Introduced Biological Control Agents. *Science*. 2001;293(5533):1314-6.

32. Munro V, Henderson I. Nontarget effect of entomophagous biocontrol: shared parasitism between native lepidopteran parasitoids and the biocontrol agent *Trigonospila brevifacies* (Diptera: Tachinidae) in forest habitats. *Environmental Entomology*. 2002;31(2):388-96.
33. Carvalheiro LG, Buckley YM, Ventim R, Fowler SV, Memmott J. Apparent competition can compromise the safety of highly specific biocontrol agents. *Ecol Lett*. 2008;11(7):690-700.
34. Veldtman R, Lado TF, Botes A, Procheş Ş, Timm AE, Geertsema H, et al. Creating novel food webs on introduced Australian acacias: indirect effects of galling biological control agents. *Diversity and Distributions*. 2011;17(5):958-67.
35. López-Núñez FA, Heleno RH, Ribeiro S, Marchante H, Marchante E. Four-trophic level food webs reveal the cascading impacts of an invasive plant targeted for biocontrol. *Ecology*. 2017;98(3):782-93.
36. Macfadyen S, Gibson R, Polaszek A, Morris RJ, Craze PG, Planqué R, et al. Do differences in food web structure between organic and conventional farms affect the ecosystem service of pest control? *Ecol Lett*. 2009;12(3):229-38.
37. Gravel D, Poisot T, Albouy C, Velez L, Mouillot D. Inferring food web structure from predator-prey body size relationships. *Methods in Ecology and Evolution*. 2013;4(11):1083-90.
38. Bartomeus I, Gravel D, Tylianakis JM, Aizen MA, Dickie IA, Bernard-Verdier M. A common framework for identifying linkage rules across different types of interactions. *Functional Ecology*. 2016;30(12):1894-903.
39. Brousseau PM, Gravel D, Handa IT. Trait matching and phylogeny as predictors of predator-prey interactions involving ground beetles. *Functional Ecology*. 2018;32(1):192-202.
40. Gravel D, Baiser B, Dunne JA, Kopelke J-P, Martinez ND, Nyman T, et al. Bringing Elton and Grinnell together: a quantitative framework to represent the biogeography of ecological interaction networks. *Ecography*. 2019;42:401-15.
41. Eklöf A, Helmus MR, Moore M, Allesina S. Relevance of evolutionary history for food web structure. *Proc R Soc B-Biol Sci*. 2012;279:1588-96.
42. Peralta G. Merging evolutionary history into species interaction networks. *Functional Ecology*. 2016;30(12):1917-25.
43. Eklöf A, Stouffer DB. The phylogenetic component of food web structure and intervality. *Theoretical ecology*. 2016;9(1):107-15.
44. Rezende EL, Lavabre JE, Guimaraes PR, Jordano P, Bascompte J. Non-random coextinctions in phylogenetically structured mutualistic networks. *Nature*. 2007;448(7156):925-8.
45. Eklöf A, Jacob U, Kopp J, Bosch J, Castro-Urgal R, Chacoff NP, et al. The dimensionality of ecological networks. *Ecol Lett*. 2013;16(5):577-83.
46. Pearse IS, Altermatt F. Predicting novel trophic interactions in a non-native world. *Ecol Lett*. 2013;16(8):1088-94.
47. Desjardins-Proulx P, Laigle I, Poisot T, Gravel D. Ecological interactions and the Netflix problem. *PeerJ*. 2017;5(8).
48. Pichler M, Boreux V, Klein AM, Schleuning M, Hartig F. Machine learning algorithms to infer trait-matching and predict species interactions in ecological networks. *Methods in Ecology and Evolution*. 2019;11:281– 93.
49. Stang M, Klinkhamer PG, van der Meijden E. Asymmetric specialization and extinction risk in plant-flower visitor webs: a matter of morphology or abundance? *Oecologia*. 2007;151(3):442-53.
50. García D, Martínez D, Stouffer DB, Tylianakis JM. Exotic birds increase generalization and compensate for native bird decline in plant-frugivore assemblages. *J Anim Ecol*. 2014;83(6):1441-50.
51. Staniczenko PPA, Lewis OT, Tylianakis JM, Albrecht M, Coudrain V, Klein AM, et al. Predicting the effect of habitat modification on networks of interacting species. *Nature Communications*. 2017;8(792).
52. Vázquez DP, Melián CJ, Williams NM, Blüthgen N, Krasnov BR, Poulin R. Species abundance and asymmetric interaction strength in ecological networks. *Oikos*. 2007;116(7):1120-7.

53. Peralta G, Perry GLW, Vázquez DP, Dehling DM, Tylianakis JM. Strength of niche processes for species interactions is lower for generalists and exotic species. *J Anim Ecol.* 2020;00:1-11.
54. Coux C, Donoso I, Tylianakis JM, García D, Martínez D, Dehling DM, et al. Tricky partners: native plants show stronger interaction preferences than their exotic counterparts. *Ecology.* 2020:e03239.
55. Laliberté E, Tylianakis JM. Deforestation homogenizes tropical parasitoid-host networks. *Ecology.* 2010;91(6):1740-7.
56. Peralta G, Frost CM, Rand TA, Didham RK, Tylianakis JM. Complementarity and redundancy of interactions enhance attack rates and spatial stability in host-parasitoid food webs. *Ecology.* 2014;95(7):1888-96.
57. Huddleston T. The braconid genus *Meteorus* in New-Zealand (Insecta: Hymenoptera). *J Nat Hist.* 1986;20(2):255-65.
58. Austin AD, Dangerfield PC. Synopsis of Australasian Microgastrinae (Hymenoptera: Braconidae), with a key to genera and description of new taxa. *Invertebrate Systematics.* 1992;6:1-76.
59. Berry JA. *Meteorus pulchricornis* (Wesmael) (Hymenoptera: Braconidae: Euphorinae), a new record for New Zealand. *New Zealand Entomologist.* 1997;20(1):45-8.
60. van Achterberg C, Berndt L, Bockerhoff E, Berry J. A new species of genus *Aleiodes* Wesmael from New Zealand (Hymenoptera: Braconidae: Rogadinae). *Zoologische Mededelingen.* 2004;78(301-311).
61. Woodward G, Ebenman B, Emmerson M, Montoya JM, Olesen JM, Valido A, et al. Body size in ecological networks. *Trends Ecol Evol.* 2005;20(7):402-9.
62. Emer C, Memmott J, Vaughan IP, Montoya D, Tylianakis JM. Species roles in plant–pollinator communities are conserved across native and alien ranges. *Diversity and Distributions.* 2016;22(8):841-52.
63. Krishna A, Guimarães PR, Jordano P, Bascompte J. A Neutral-Niche Theory of Nestedness in Mutualistic Networks. *Oikos.* 2008;117:1609-18.
64. Vázquez DP, Chacoff NP, Cagnolo L. Evaluating Multiple Determinants of the Structure of Plant-Animal Mutualistic Networks. *Ecology.* 2009;90(8):2039-46.
65. Albrecht J, Bohle V, Berens DG, Jaroszewicz B, Selva N, Farwig N. Variation in neighbourhood context shapes frugivore-mediated facilitation and competition among co-dispersed plant species. *J Ecol.* 2015;103(2):526-36.
66. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in Python. *Journal of machine learning research.* 2011;12:2825-30.
67. Dormann CF, Fründ J, Blüthgen N, Gruber B. Indices, graphs and null models: analyzing bipartite ecological networks. *The Open Ecology Journal.* 2009;2:7-24.
68. Moran MD. Arguments for rejecting the sequential Bonferroni in ecological studies. *Oikos.* 2003;100:403-5.
69. Nakagawa S. A farewell to Bonferroni: the problems of low statistical power and publication bias. *Behavioral Ecology.* 2004;15(6):1044-5.
70. Zuur A, Ieno EN, Walker N, Saveliev AA, Smith GM. *Mixed Effects Models and Extensions in Ecology with R*: Springer; 2009.
71. Barton K. MuMIn: Multi-Model Inference. R package, <https://CRAN.R-project.org/package=MuMIn2020>.
72. Burnham KP, Anderson DR. P values are only an index to evidence: 20th-vs. 21st-century statistical science. *Ecology.* 2014;95(3):627-30.
73. RCoreTeam. R: A language and environment for statistical computing (R Foundation for Statistical Computing); URL <https://www.R-project.org/>. 2018.
74. Bates D, Mächler M, Bolker B, Walker S. Fitting Linear Mixed-Effects Models using lme4. *Journal of Statistical Software.* 2015;67(1):1-48.
75. Lefcheck JS. piecewiseSEM: Piecewise structural equation modelling in R for ecology, evolution, and systematics. *Methods in Ecology and Evolution.* 2016;7(5):573-9.

76. Harrison XA. Using observation-level random effects to model overdispersion in count data in ecology and evolution. *PeerJ*. 2014;2:e616.
77. Bolker BM, Brooks ME, Clark CJ, Geange SW, Poulsen JR, Stevens MHH, et al. Generalized linear mixed models: a practical guide for ecology and evolution. *Trends Ecol Evol*. 2009;24(3):127-35.
78. Crawley MJ. *Statistics: an introduction using R*. Second ed. Chichester, West Sussex: Wiley; 2014.
79. Louda SM, Pemberton RW, Johnson MT, Follett PA. Nontarget effects-the achilles' heel of biological control? Retrospective Analyses to Reduce Risk Associated with Biocontrol Introductions. *Annual Review of Entomology*. 2003;48(1):365-96.
80. Matsubayashi KW, Ohshima I, Nosil P. Ecological speciation in phytophagous insects. *Entomologia Experimentalis et Applicata*. 2010;134(1):1-27.
81. Berlocher SH, Feder JL. Sympatric speciation in phytophagous insects: moving beyond controversy? *Annual Review of Entomology*. 2002;47(1):773-815.
82. Wood T, Tilmon K, Shantz A, Harris C, Pesek J. The role of host-plant fidelity in initiating insect race formation. *Evolutionary Ecology Research*. 1999;1(3):317-32.
83. Bush GL. Modes of animal speciation. *Annu Rev Ecol Syst*. 1975;6(1):339-64.
84. Stireman JO, Nason JD, Heard SB, Seehawer JM. Cascading host-associated genetic differentiation in parasitoids of phytophagous insects. *Proceedings of the Royal Society of London B: Biological Sciences*. 2006;273(1586):523-30.
85. Hamerlinck G, Hulbert D, Hood GR, Smith JJ, Forbes AA. Histories of host shifts and cospeciation among free-living parasitoids of Rhagoletis flies. *Journal of Evolutionary Biology*. 2016;29(9):1766-79.
86. Feder JL, Forbes AA. Sequential speciation and the diversity of parasitic insects. *Ecological Entomology*. 2010;35:67-76.
87. Hood GR, Forbes AA, Powell TH, Egan SP, Hamerlinck G, Smith JJ, et al. Sequential divergence and the multiplicative origin of community diversity. *Proceedings of the National Academy of Sciences*. 2015;112(44):E5980-E9.
88. Kankare M, van Nouhuys S, Hanski I. Genetic Divergence Among Host-Specific Cryptic Species in *Cotesia melitaearum* Aggregate (Hymenoptera: Braconidae), Parasitoids of Checkerspot Butterflies. *Annals of the Entomological Society of America*. 2005;98(3):382-94.
89. Jager CR, Menken B. Sympatric speciation in *Ageniaspis fuscicollis* (Hymenoptera: Encyrtidae), and endoparasitoid of *Yponomeuta* spp. (Lepidoptera: Yponomeutidae), revealed by RAPD markers. *Proc Exp Appl Entomol*. 1994;5:19-24.
90. Henri DC, van Veen FF. Body size, life history and the structure of host-parasitoid networks. *Advances in Ecological Research*. 2011;45:135-80.
91. Tylianakis JM, Tscharntke T, Lewis OT. Habitat modification alters the structure of tropical host-parasitoid food webs. *Nature*. 2007;445(7124):202-5.
92. Klein A-M, Steffan-Dewenter I, Tscharntke T. Rain Forest Promotes Trophic Interactions and Diversity of Trap-Nesting Hymenoptera in Adjacent Agroforestry. *J Anim Ecol*. 2006;75(2):315-23.
93. Peralta G, Vázquez DP, Chacoff NP, Lomáscolo SB, Perry GL, Tylianakis JM. Trait matching and phenological overlap increase the spatio-temporal stability and functionality of plant-pollinator interactions. *Ecol Lett*. 2020.
94. Peralta G, Frost CM, Didham RK. Plant, herbivore and parasitoid community composition in native Nothofagaceae forests vs. exotic pine plantations. *Journal of applied ecology*. 2018;55(3):1265-75.
95. Peralta G, Frost CM, Didham RK, Rand TA, Tylianakis JM. Non-random food-web assembly at habitat edges increases connectivity and functional redundancy. *Ecology*. 2017;98(4):995-1005.
96. Jordano P. Sampling networks of ecological interactions. *Functional Ecology*. 2016;30(12):1883-93.

97. Jordano P. Patterns of mutualistic interactions in pollination and seed dispersal: connectance, dependence asymmetries, and coevolution. *The American Naturalist*. 1987;129(5):657-77.
98. Olesen JM, Bascompte J, Dupont YL, Elberling H, Rasmussen C, Jordano P. Missing and forbidden links in mutualistic networks. *Proceedings of the Royal Society B: Biological Sciences*. 2011;278(1706):725–32.
99. Aizen MA, Morales CL, Morales JM. Invasive mutualists erode native pollination webs. *PLoS Biology*. 2008;6(2):e31.
100. Lopezaraiza-Mikel ME, Hayes RB, Whalley MR, Memmott J. The impact of an alien plant on a native plant-pollinator network: an experimental approach. *Ecol Lett*. 2007;10(7):539-50.

Supporting information

Appendix S1. Trait and phylogenetic details, machine-learning model details, hyperparameter

tuning details, supplementary analysis and coefficient tables for analyses.

Appendix S2. Host-parasitoid interaction data (test sites). Note: N= native forest, P=plantation

forest, EH=host reduction treatment, EC=control, extra=whether the caterpillar was collected during

extra sampling, B=before host-reduction treatment (i.e. at time t), A=after host-reduction treatment

(i.e. at time $t+7$)

Appendix S3. Host body size data.

Appendix S4. Parasitoid body size data.

Appendix S5. GenBank accession numbers for host species.

Appendix S6. GenBank accession numbers for parasitoid species.

Appendix S7. Host-parasitoid interaction data (training sites).

A **B**