

Statistical Report

Ecology, 99(6), 2018, pp. 1277–1283
© 2018 by the Ecological Society of America

Unravelling changing interspecific interactions across environmental gradients using Markov random fields

NICHOLAS J. CLARK ^{1,4}, KONSTANS WELLS ², AND OSCAR LINDBERG³

¹*School of Veterinary Science, University of Queensland, Gatton, Queensland 4343 Australia*

²*Environmental Futures Research Institute, School of Environment, Griffith University, Brisbane, Queensland 4111 Australia*

³*Department of Mathematics and Statistics, University of Turku, 20500 Turku Finland*

Abstract. Inferring interactions between co-occurring species is key to identify processes governing community assembly. Incorporating interspecific interactions in predictive models is common in ecology, yet most methods do not adequately account for indirect interactions (where an interaction between two species is masked by their shared interactions with a third) and assume interactions do not vary along environmental gradients. Markov random fields (MRF) overcome these limitations by estimating interspecific interactions, while controlling for indirect interactions, from multispecies occurrence data. We illustrate the utility of MRFs for ecologists interested in interspecific interactions, and demonstrate how covariates can be included (a set of models known as Conditional Random Fields, CRF) to infer how interactions vary along environmental gradients. We apply CRFs to two data sets of presence–absence data. The first illustrates how blood parasite (*Haemoproteus*, *Plasmodium*, and nematode microfilaria spp.) co-infection probabilities covary with relative abundance of their avian hosts. The second shows that co-occurrences between mosquito larvae and predatory insects vary along water temperature gradients. Other applications are discussed, including the potential to identify replacement or shifting impacts of highly connected species along climate or land-use gradients. We provide tools for building CRFs and plotting/interpreting results as an R package.

Key words: *co-infection; environmental gradient; graphical network model; Haemoproteus; interspecific interactions; Markov random fields; network modeling; Plasmodium; species distribution model.*

INTRODUCTION

Interactions between co-occurring species are crucial for explaining variation in biodiversity (Dayton 1971, Azaele et al. 2010, Mayfield and Stouffer 2017). Understanding these interactions is necessary to predict how communities will respond to climate change or habitat modification (Götzenberger et al. 2012, Golding and Harris 2013, Harris 2015a). A multitude of statistical approaches can model species' pairwise co-occurrence probabilities, typically by estimating interspecific associations as multivariate variance–covariance matrices after controlling for species' environmental niches (Golding et al. 2015, Tikhonov et al. 2017). These so-called Joint Species Distribution Models (JSDMs) are used widely in ecology and have yielded important insights, ranging from improved species distribution models to the discovery of unforeseen parasite co-infection patterns (Kissling et al. 2012, Clark et al. 2016, Ovaskainen et al. 2017).

JSDMs open exciting avenues for predicting assemblage variation while accounting for biotic interactions. However, most JSDMs are unable to incorporate non-stationarity in pairwise dependencies (e.g., variation in species interactions

across spatiotemporal gradients; Kissling et al. 2012, but see Tikhonov et al. 2017). The assumption that interactions are constant limits predictive capacity, as both the direction and strength of interspecific interactions can vary across environmental gradients (Grime 1973, He et al. 2013). Interpretation of covariance coefficients as interspecific “interactions” has also been questioned, as pairwise interactions are not directly accounted for (they are inferred from correlated residuals), nor are their effect sizes directly comparable to those of covariates (Pellissier et al. 2013, Harris 2015a, Tikhonov et al. 2017).

Graphical network models offer the opportunity to flexibly estimate conditional dependencies from networks of interacting variables (Meinshausen and Bühlmann 2006, Azaele et al. 2010). Markov random fields (MRFs; also known as Markov networks) are particularly useful graphical network models as they estimate dependencies among presence–absence variables (commonly used in species distribution modeling) and their properties are well understood (Ising 1925, Li et al. 2012). Because MRFs are undirected (edges do not incorporate directionality), cyclicity is feasible. This means we can ask if each variable pair is conditionally dependent given their relationships with all other variables, as opposed to estimating independence given only relationships of their parent variables (as in directed Bayesian networks). In other words, MRFs infer direct interactions (i.e.,

Manuscript received 16 January 2018; revised 20 February 2018; accepted 2 March 2018. Corresponding Editor: Caz M. Taylor.

⁴E-mail: nicholas.j.clark1214@gmail.com

competition, facilitation) that may be obscured by strong indirect interactions (i.e., competitors that both show strong correlations with a third species; see Harris 2016: Fig. 1 for a visualization of direct vs. indirect interactions). MRFs have been implemented across many disciplines, with aims ranging from identifying gene associations and mutation dependencies to uncovering network relationships in language passages (Li et al. 2012, Lee and Hastie 2015). Recently, their flexibility has been expanded by incorporating covariates to estimate conditional dependencies among variables and how these dependencies vary with covariates (also known as conditional random fields, or CRFs; Sutton and McCallum 2012, Cheng et al. 2014, Lindberg 2016).

Although less often used in ecology (but see Azaele et al. 2010, Harris 2016) graphical network models provide promising frameworks to understand interactions and their consequences on biodiversity. Here, we adapt a CRF model described by Cheng et al. (2014) to identify co-occurrence probabilities and estimate how they vary across environmental gradients. We apply models to two community data sets and, using functions in the MRFCov R package (Clark et al. 2018), we demonstrate that model execution and interpretation are straightforward. Our aim is not to carry out a statistical proof of the model (provided by works cited in this paper, most notably by Cheng et al. 2014) but to facilitate application of CRFs for analyzing ecological data.

METHODS AND CASE STUDIES

Approximating conditional random fields with logistic regressions

A CRF models all the species jointly in a unified graphical model. However, we can focus on one species at a time within the larger model to better understand how pairwise interactions are modeled conditionally on an external covariate. The log-odds of observing species j given the presence-absence of species k can be modeled with inclusion of covariate x by:

$$\log\left(\frac{P(y_j = 1|y_{\setminus j}, x)}{1 - P(y_j = 1|y_{\setminus j}, x)}\right) = \alpha_{j0} + \beta_j^T x + \sum_{k:k \neq j} (\alpha_{jk0} + \beta_{jk}^T x) y_k \quad (1)$$

where \mathbf{y}_j is a vector of binary observations for species j (1 if the species was observed, 0 if not), $\mathbf{y}_{\setminus j}$ is a vector of binary observations for all other species apart from j , α_{j0} is the species-level intercept and β_j^T is the coefficient of covariate x on species j 's occurrence probability. Linear parameterization of the CRF [$\alpha_{j0} + \beta_j^T x + \sum_{k:k \neq j} (\alpha_{jk0} + \beta_{jk}^T x) y_k$] can be estimated using logistic regression, where each coefficient represents the size of a predictor's effect on a species' conditional log-odds. For instance, if $\alpha_{jk0} = 0$, occurrence probabilities of species j and k are conditionally independent, after accounting for covariates and other species. If $\alpha_{jk0} \neq 0$ but $\beta_{jk}^T x = 0$, the species' occurrence probabilities are conditionally dependent, but the strength of this dependence does not vary with covariate x . For a data set with J species and X covariates, maximizing the joint likelihood requires

estimation of $(J + 1) \times X(X + 1)/2$ coefficients in addition to a normalizing constant that grows exponentially with the number of species (Cheng et al. 2014). This quickly becomes intractable even for moderately sized data sets. A common approach to circumvent the exponentially growing parameter estimation is to approximate the graph by employing a series of single-species regressions and combining the parameter vectors in a common matrix (Cheng et al. 2014, Harris 2016). Because CRFs are undirected, coefficients of conditional dependence must be symmetric, where $\alpha_{jk0} = \alpha_{kj0}$. Likewise, coefficients for covariate effects on co-occurrences must be symmetric ($\beta_{jk}^T x = \beta_{kj}^T x$). As these estimates come from separate regressions, symmetry is not guaranteed. We follow Cheng et al. (2014) to enforce symmetry using a *separate.max* approach to keep the coefficient with the larger absolute value. In this way, we generate a unified graphical network by separately maximizing the conditional log-likelihood of each species. Studies with many species or covariates may benefit from a conservative *separate.min* approach, keeping the coefficient with the smaller absolute value.

As with JSDMs, CRFs can be limited by the complexity of estimating dependency parameters from data sets with many species and/or covariates (Golding et al. 2015, Harris 2015b). For example, the mosquito data set described in Case study 2 includes 16 species and 17 covariates, resulting in 288 coefficients in each regression. As estimating so many coefficients can lead to problems of overfitting, sparsity is added using regularization, which forces some coefficients to zero (Goeman 2010). Regularization is used in a variety of statistical frameworks due to its ability to simultaneously perform variable selection (the least important variables are more likely to be penalized at lower regularization values) and model shrinkage to avoid overfitting. We use L_1 regularization (e.g., LASSO), where lower values of the regularization parameter λ lead to a greater number of non-zero coefficients. Note that regularization is influenced by the scale of predictors (covariates on larger scales are less likely to be penalized, even if they have marginal effects; Goeman 2010) so it is important that covariates are appropriately scaled. Following Lindberg (2016), we use 10-fold cross validation to identify optimal λ values. Considering that separate regressions are unified in a joint model, we estimate model fit by comparing predicted and observed outcomes simultaneously for all species. Cross validation is preferred to statistics such as AIC or BIC, which are difficult to interpret in LASSO applications (Goeman 2010).

Case study 1: Estimating co-infection probabilities for avian blood parasites

Data published by Clark et al. (2016) contains binary observations for four vector-borne blood parasites (*Haemoproteus zosteropsis* species complex, *H. killangoi* species complex, *Plasmodium* spp., and nematode microfilaria species) among New Caledonian birds. The authors used a JSDM to estimate co-infection probabilities in the primary host group (species in the genus *Zosterops*; $n = 449$), including a covariate representing the proportion of captures in each site attributed to *Zosterops* spp. (a proxy of host relative abundance). We did not impute missing microfilaria data as in

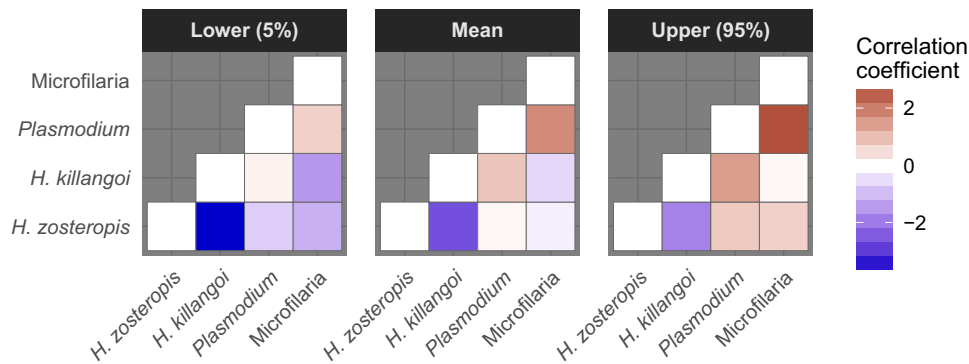
the original analysis (some birds were not tested for microfilaria), instead treating missing data as zeros. We first estimated an MRF (no covariates) and compared to a model fit using the *rosalia* R package (Harris 2015b), which solves small MRF networks by maximum likelihood (Harris 2016). Interaction signs showed perfect concordance (both models estimated the same positive and negative interactions), suggesting our approach does well at approximating the MRF.

We next estimated a CRF by including *Zosterops* relative abundance as a scaled continuous predictor (mean centered, divided by 1 SD). To choose a regularization parameter, we tested predictive performance at a range of λ s (values between 1.0 and 2.5, sequenced in 0.1 steps) using 10-fold cross validation. We fit models at each λ using a training partition containing ~90% of observations ($n \sim 404$), and calculated linear predictions of species' occurrences for a test partition containing ~10% of observations ($n \sim 45$; Appendix S1). This was repeated 10 times (training partitions chosen at random) and, for each test prediction, we calculated sensitivity (the proportion of observed positives that were predicted to be positive). To assess whether addition of the covariate improved predictions, we also assessed performance of MRFs (without a covariate) at each λ . Using LOESS regressions to analyze sensitivity vs. λ values, we find that (1) inclusion of the covariate increases model sensitivity across the range of λ s and (2) increased

regularization decreases our ability to predict true positives (Appendix S2: Fig. S1). Given our interest in predicting relatively rare parasite infections, fitting models at the λ that coincides with the asymptote for sensitivity ($\lambda = 1.5$; Appendix S2: Fig. S1) is a reasonable approach (higher λ s lead to less chance of overfitting). Yet lower λ s will not always overfit, and may provide better insights if higher λ s regularize some interactions. We chose to fit models using λ s within the range uniform[1, 1.5]. For each λ , we modeled 500 bootstrapped replicates (randomly shuffling rows of training sets with replacement) to account for uncertainty in interaction coefficients.

Coefficient 95% confidence intervals identified several strongly supported interactions (Fig. 1a). Microfilaria and *Plasmodium* spp. showed a positive interaction, while the two *Haemoproteus* species showed a strong negative interaction (Fig. 1a). We compared interactions to those from competing methods to validate our approach (R scripts available in Data S1). We fitted a Bayesian JSDM (using the R package BayesComm; Golding and Harris 2013) and estimated species' pairwise correlations and partial correlations following methods in Harris (2016). Our approach identified similar interactions to those from Clark et al. (2016) and to competing methods, with methods that account for environment (BayesComm and environmentally controlled partial correlations) showing the most similar

A) Predicted interactions (95% CI)



B) Predicted interactions at varying *Zosterops* relative abundance

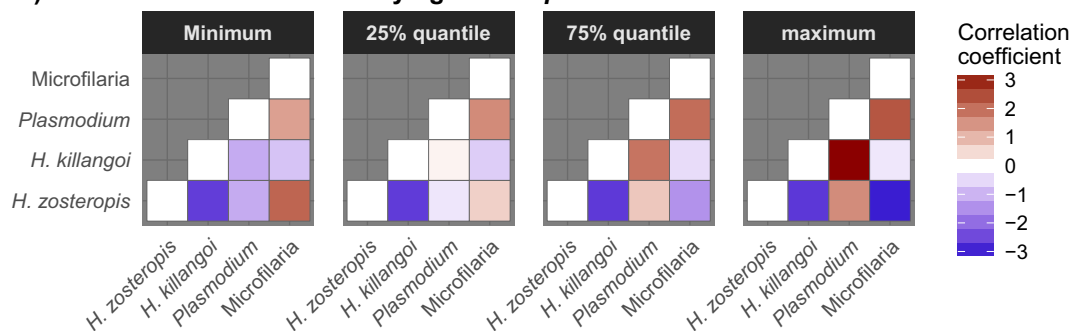


FIG. 1. Predicted interactions between blood parasites (*Haemoproteus zosteropsis* species complex, *H. killangoi* species complex, *Plasmodium* spp. and nematode microfilaria species). Positive interactions are in red; negative interactions are in blue. Panel A shows 95% confidence intervals of coefficients estimated from bootstrapped conditional random fields (CRF) models (including *Zosterops* relative abundance as a covariate) fit in the optimal range of λ values. Note that these estimates represent predicted interactions when the scaled covariate is zero, representing mean *Zosterops* spp. relative abundance. Panel B shows how interactions are predicted to change as *Zosterops* spp. relative abundance changes.

effects to the CRF (Appendix S2: Fig. S2). However, competing methods provide little information on species' co-occurrences beyond these average interactions. Our approach, in contrast, provides two powerful new insights: (1) perhaps representing the most attractive property of CRFs, we can directly compare interaction coefficients to those of other predictors. Interspecific interactions were stronger predictors than was the host relative abundance covariate for each of the four parasites (see Table 1 for the *H. zosteropsis* model summary; Appendix S1). By contrast, comparing effects of interactions to those of other predictors is not straightforward with most JSDMs (typically requiring variance partitioning among fixed and random effects; see Ovaskainen et al. 2017); (2) CRFs revealed that several key interactions were predicted to change as host relative abundance changes (Fig. 1b; Appendix S1). Most competing methods estimate static interactions, providing no indication of how interactions change (but see Tikhonov et al. [2017], who relate correlations to covariates, though the ability to compare interaction effect sizes is still difficult). For both *Haemoproteus* spp., interactions with *Plasmodium* spp. became more positive with increasing host relative abundance, suggesting these parasites are more likely to co-occur in sites with high host availability, while the interaction between *H. zosteropsis* and microfilaria showed the opposite (Fig. 1b). The negative interaction between *Haemoproteus* spp. did not vary.

We explored how well our model predicted variable interactions by subsetting the data to only include sites with above average *Zosterops* relative abundance. We then estimated environmentally controlled partial correlations and compared these to interactions predicted from our model. This was repeated using sites with lower than average host abundance (R scripts available in Data S1). Five of the six possible pairwise interactions were successfully predicted at low host abundance (based on overlapping confidence intervals), while only four were predicted at high host abundance (Fig. 2). All four parasites were less likely to occur in sites with high host abundance, which may help to explain the wider partial correlations confidence intervals and the poorer matching with CRF predictions (Fig. 2). Overall, our findings raise questions about the roles that host availability and parasite community

composition play in driving parasite co-occurrences (Clark et al. 2016).

Case study 2: Estimating biotic interactions influencing larval mosquito occurrences

Data published by Golding et al. (2015) contains binary observations for four mosquito species and 12 other benthic faunal groups recorded across three UK marsh sites. The data provides four continuous covariates (depth, salinity, oxidation reduction potential, and temperature) and binary covariates for a range of vegetation indicators. The authors used stepwise selection to determine important abiotic covariates for each species; they then used a JSDM to identify co-occurrence probabilities after accounting for specified covariates in each species-level regression. The larger size of the mosquito data set (16 species and 17 predictors) affords an opportunity to test our approach's performance in larger parameter spaces. As above, we estimated performance using 10-fold cross validation of bootstrapped data sets. Considering the high number of species and their variable occurrence frequencies, our criterion for choosing λ s was to identify a range of values that maintained high proportions of true predictions ($\lambda = \text{uniform}[4.8, 5.8]$; Appendix S3).

Models identified numerous interspecific interactions. We found a strong positive interaction between *Culex* mosquito species (*modestus* and *pipiens*; Appendix S3), suggesting that they rely on similar larval habitats. We identified negative ditch shrimp (*Palaemonetes* spp.)–*Culex modestus*, and saucer bug–*Anopholes maculipennis* interactions, indicating possible predatory exclusion (Appendix S3). Our model estimated that several interactions between mosquito species and potential predators changed considerably across temperature gradients. Saucer bugs (*Ilyocoris* spp.) were less likely to co-occur with *A. maculipennis* as temperature increased, while the damselflies (*Coleoptera* spp.)–*Culex modestus* interaction showed the opposite (Fig. 3). Several interactions also changed across gradients of depth and oxidation reduction potential (Appendix S3). Considering the roles of these mosquito species in disease transmission, identifying variable interactions across environmental gradients can be informative for vector management programs.

DISCUSSION

In most ecological applications, the occurrence probabilities of species, and the way species interact within a community, cannot be assumed to be static or independent of environmental influences. Given the unprecedented environmental change occurring globally, estimating variable species interactions across temporal or environmental gradients is of vital importance (Wilson and Keddy 1986). We suggest CRFs can help achieve this goal. Applying CRFs to two case studies, we build on inferences that can be gleaned from JSDMs to show that interactions between avian blood parasites have strong influences on infection rates and that co-infection probabilities are not static across host abundance gradients. In addition, we found that potential interactions between larvae of disease-transmitting mosquitoes and their

TABLE 1. Example model summary showing mean regression coefficients (β) and their relative importance for predictors of *Haemoproteus zosteropsis* occurrence, taken across bootstrapped CRF models (using λ s in the range [1, 1.5]).

| Predictor | Mean coefficient (β) | Relative importance ($\beta/\text{sum}(\beta^2)$) |
|---|------------------------------|---|
| <i>H. killangoi</i> | –2.69 | 0.62 |
| Host relative abundance \times Microfilaria spp. | –1.30 | 0.14 |
| Microfilaria spp. | 1.06 | 0.10 |
| Host relative abundance | –0.94 | 0.08 |
| Host relative abundance \times <i>Plasmodium</i> spp. | 0.70 | 0.04 |
| <i>Plasmodium</i> spp. | –0.55 | 0.03 |

Note: Predictors with multiplier symbol (\times) indicate an interspecific interaction that varies with varying host relative abundance.

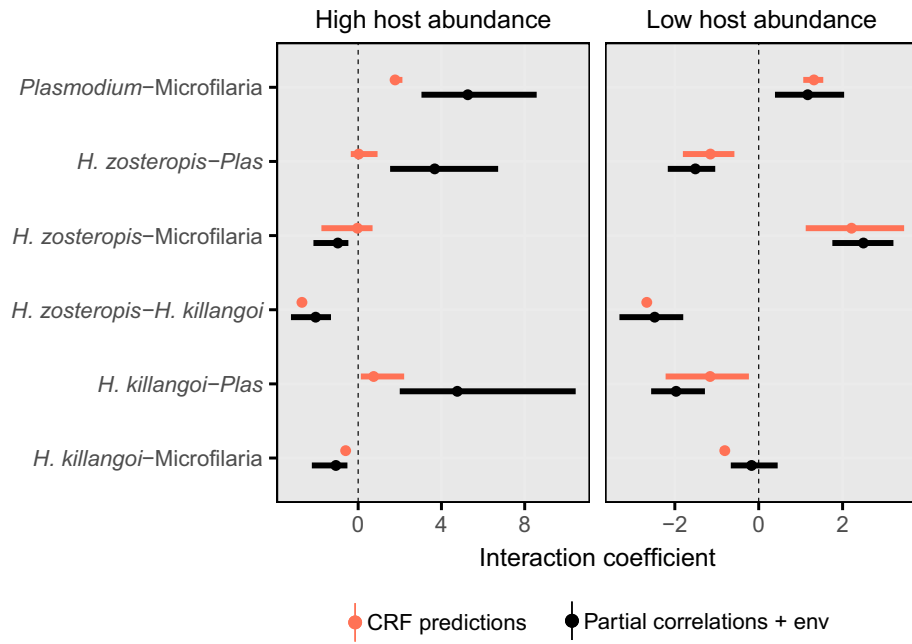


FIG. 2. CRF predicted interactions for subsets of the bird parasite data set. Predictions are compared to partial correlations that represent estimates of partial correlations in species' occurrences after controlling for correlations with remaining species and correlations with environmental covariates (env). The left panel shows estimates for sites with above average *Zosterops* spp. relative abundance; the right panel shows estimates for sites with below average relative abundance. Points show estimate medians; bars show 95% confidence intervals from 1,000 bootstrapped estimates.

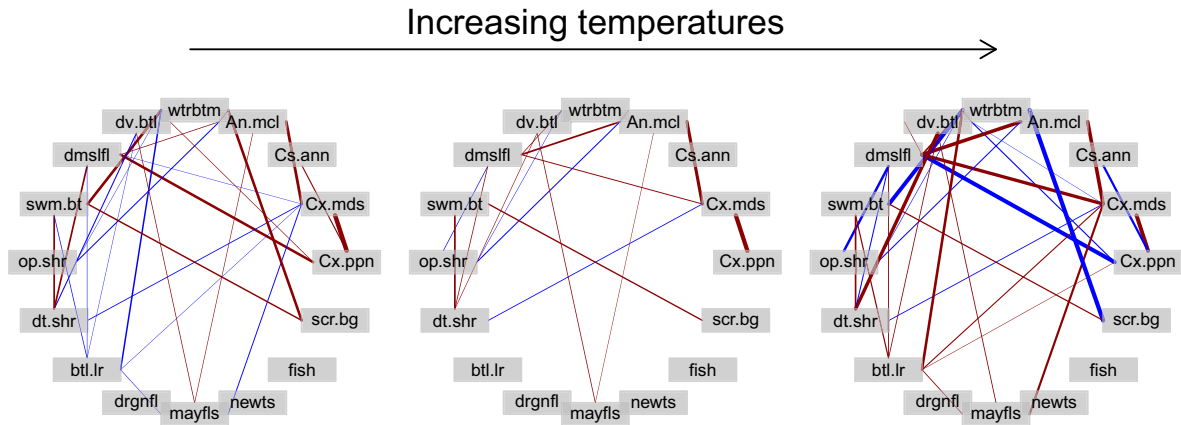


FIG. 3. Predicted variation in interactions between mosquito species and other marsh-dwelling species across water temperature gradients. Positive interactions are in red; negative interactions in blue. Line thicknesses indicate the relative strength of interactions. An.mcl, *Anopheles maculipennis*; Cs.ann, *Culiseta annulata*; Cx.mds, *Culex modestus*; Cx.ppn, *Culex pipiens*; scr.bg, saucer bugs (Ilyocoris); fish, Pisces; newts, Pleurodelinae; mayfls, mayflies (Ephemeroptera); drgnfl, dragonflies (Anisoptera); btl.lr, beetle larvae (Coleoptera); dt.shr, ditch shrimp (Palaemonetes); op.shr, amphipods (Gammaridae); swm.bt, swimming beetles (Haliplidae); dmslfl, damselflies (Zygoptera); dv.btl, diving beetles (Dysticidae); wtrbtm, waterboatmen (Corixidae).

co-occurring predators may vary across gradients of water depth, temperature, and oxidation reduction potential.

Our report offers a glimpse of the promising ways that CRFs can infer multispecies interactions. The ability to include covariates, directly estimate interaction terms, and evaluate relative effect sizes provides a framework to extend JSDMs to multilayer ecological networks (Pilosof et al. 2017) and opens a wealth of useful directions. An attractive property of CRFs is their ability to identify highly connected variables with particularly large influences on graph structure. Applied to ecological settings, uncovering such highly

connected species can provide evidence of species whose presences have important consequences for community composition. Whether highly connected species are ecologically replaced across gradients of land use, elevation, or temperature, or whether their impact on other species changes in more diverse assemblages, are topical questions (Tanner et al. 1994, Le Roux et al. 2005) that can be addressed using CRFs.

By presenting analyses that can be programmed using an open-source R package (Clark et al. 2018), this paper extends the relatively unfamiliar topic of CRFs (and graphical networks in general) to the applied ecologist's analysis

toolset. Considering the key roles biotic interactions play in generating biodiversity, influencing species' distributions, and driving infection epidemics (Kissling et al. 2012, He et al. 2013), there is need for analytical approaches that generate interpretable estimates of interactions from types of data commonly gathered by ecologists. In this respect, coefficients estimated from CRFs will be of broad interest because they (1) offer the opportunity to disentangle complex biotic interactions from environmental drivers and (2) are generated from arguably one of the most common analysis procedures in ecology: logistic regression of presence-absence data. Interpreting CRF coefficients will be straightforward for a broad range of ecologists working in different fields. Moreover, it is important to note that, while our study describes models built with presence-absence data, the ability to include covariates in interspecific interaction studies has also been introduced as an extension to Gaussian graphical models to model continuous outcomes (i.e., abundances or fitness measures; Li et al. 2012). Additionally, methods implementing mixed graphical models, which can incorporate both binary and continuous data, are being investigated to answer a growing need to model relationships between variables of heterogeneous variable types (Lee and Hastie 2015). These will add further flexibility to the types of interactions that can be explored with undirected graphical models.

We caution that our method is not a replacement for hierarchical JSDMs. The ability to account for species' non-linear responses to environmental covariates in distribution models affords greater flexibility for forecasting regime changes (Harris 2015a). While our approach will be useful for accurately estimating interaction networks from moderate- or large-sized data sets, independent regressions also prevent the use of informative grouping parameters that can be modeled with hyperpriors in multispecies models (Ovaskainen et al. 2017). Although we have much to learn about species interactions and their effects on biodiversity, we hope that CRFs will serve as a useful addition in the expanding toolbox for understanding community assembly.

ACKNOWLEDGMENTS

We thank J. Giles and S. Clegg for discussions that stimulated the research. The work was partially funded by a National Geographic Society Grant (#9383-13). David J Harris and an anonymous reviewer provided helpful criticisms that improved the quality of the manuscript.

LITERATURE CITED

- Azaele, S., R. Muneeppeerakul, A. Rinaldo, and I. Rodriguez-Iturbe. 2010. Inferring plant ecosystem organization from species occurrences. *Journal of Theoretical Biology* 262:323–329.
- Cheng, J., E. Levina, P. Wang, and J. Zhu. 2014. A sparse Ising model with covariates. *Biometrics* 70:943–953.
- Clark, N. J., K. Wells, D. Dimitrov, and S. M. Clegg. 2016. Co-infections and environmental conditions drive the distributions of blood parasites in wild birds. *Journal of Animal Ecology* 85:1461–1470.
- Clark, N. J., K. Wells, and O. Lindberg. 2018. MRFcov: Markov random fields with additional covariates. R package version 1.0. GitHub, <https://github.com/nicholasjclark/MRFcov>
- Dayton, P. K. 1971. Competition, disturbance, and community organization: the provision and subsequent utilization of space in a rocky intertidal community. *Ecological Monographs* 41:351–389.
- Goeman, J. J. 2010. L1 penalized estimation in the Cox proportional hazards model. *Biometrical Journal* 52:70–84.
- Golding, N., and D. J. Harris. 2013. BayesComm: Bayesian community ecology analysis. R package version 0.1-0. <https://cran.r-project.org/web/packages/BayesComm/index.html>
- Golding, N., M. A. Nunn, and B. V. Purse. 2015. Identifying biotic interactions which drive the spatial distribution of a mosquito community. *Parasites & Vectors* 8:367.
- Götzenberger, L., F. de Bello, K. A. Bråthen, J. Davison, A. Dubuis, A. Guisan, J. Lepš, R. Lindborg, M. Moora, and M. Pärtel. 2012. Ecological assembly rules in plant communities—approaches, patterns and prospects. *Biological Reviews* 87:111–127.
- Grime, J. P. 1973. Competitive exclusion in herbaceous vegetation. *Nature* 242:344–347.
- Harris, D. J. 2015a. Generating realistic assemblages with a joint species distribution model. *Methods in Ecology and Evolution* 6:465–473.
- Harris, D. J. 2015b. rosalia: exact inference for small binary Markov networks. R package version 0.1.0. Zenodo. <https://doi.org/10.5281/zenodo.17808>
- Harris, D. J. 2016. Inferring species interactions from co-occurrence data with Markov networks. *Ecology* 97:3308–3314.
- He, Q., M. D. Bertness, and A. H. Altieri. 2013. Global shifts towards positive species interactions with increasing environmental stress. *Ecology Letters* 16:695–706.
- Ising, E. 1925. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei* 31:253–258.
- Kissling, W. D., et al. 2012. Towards novel approaches to modelling biotic interactions in multispecies assemblages at large spatial extents. *Journal of Biogeography* 39:2163–2178.
- Le Roux, P. C., M. A. McGeoch, M. J. Nyakatya, and S. L. Chown. 2005. Effects of a short-term climate change experiment on a sub-Antarctic keystone plant species. *Global Change Biology* 11:1628–1639.
- Lee, J. D., and T. J. Hastie. 2015. Learning the structure of mixed graphical models. *Journal of Computational and Graphical Statistics* 24:230–253.
- Li, B., H. Chun, and H. Zhao. 2012. Sparse estimation of conditional graphical models with application to gene networks. *Journal of the American Statistical Association* 107:152–167.
- Lindberg, O. 2016. Markov random fields in cancer mutation dependencies. Thesis. University of Turku, Turku, Finland.
- Mayfield, M. M., and D. B. Stouffer. 2017. Higher-order interactions capture unexplained complexity in diverse communities. *Nature Ecology & Evolution* 1:0062.
- Meinshausen, N., and P. Bühlmann. 2006. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics* 34:1436–1462.
- Ovaskainen, O., G. Tikhonov, A. Norberg, F. Guillaume Blanchet, L. Duan, D. Dunson, T. Roslin, and N. Abrego. 2017. How to make more out of community data? A conceptual framework and its implementation as models and software. *Ecology Letters* 20:561–576.
- Pellissier, L., R. P. Rohr, C. Ndiribe, J.-N. Pradervand, N. Salamin, A. Guisan, and M. Wisz. 2013. Combining food web and species distribution models for improved community projections. *Ecology and Evolution* 3:4572–4583.
- Pilosof, S., M. A. Porter, M. Pascual, and S. Kéfi. 2017. The multi-layer nature of ecological networks. *Nature Ecology & Evolution* 1:0101.
- Sutton, C., and A. McCallum. 2012. An introduction to conditional random fields. *Foundations and Trends in Machine Learning* 4:267–373.

- Tanner, J. E., T. P. Hughes, and J. H. Connell. 1994. Species coexistence, keystone species, and succession: a sensitivity analysis. *Ecology* 75:2204–2219.
- Tikhonov, G., N. Abrego, D. Dunson, and O. Ovaskainen. 2017. Using joint species distribution models for evaluating how species-to-species associations depend on the environmental context. *Methods in Ecology and Evolution* 8:443–452.
- Wilson, S. D., and P. A. Keddy. 1986. Measuring diffuse competition along an environmental gradient: results from a shoreline plant community. *American Naturalist* 127:862–869.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at <http://onlinelibrary.wiley.com/doi/10.1002/ecy.2221/suppinfo>

DATA AVAILABILITY

Data associated with this manuscript is available from Zenodo: <https://doi.org/10.5281/zenodo.1199673>