

Uncovering bidirectional ecological associations from co-occurrence and environmental data

Learning bidirectional ecological associations [running headline]

KEYWORDS: Biotic associations, ecological networks, habitat suitability models, network inference, plant interactions, representation learning, simulated communities.

AUTHORS:

Sara Si-moussi sara.si-moussi@inrae.fr
INRAE, IRD, CIRAD, Montpellier SupAgro, Université Montpellier, UMR Eco&Sols

Esther Galbrun esther.galbrun@uef.fi
School of Computing, University of Eastern Finland

Mickaël Hedde mickael.hedde@inrae.fr
INRAE, IRD, CIRAD, Montpellier SupAgro, Université Montpellier, UMR Eco&Sols

Wilfried Thuiller wilfried.thuiller@univ-grenoble-alpes.fr
Univ. Grenoble Alpes, CNRS, Univ. Savoie Mont Blanc, CNRS, LECA, Laboratoire d'Écologie Alpine

Correspondence to **Sara Si-Moussi**

Mailing adress: 2 Place Pierre Viala 34060 Montpellier, France

E-mail: sara.si-moussi@inrae.fr *Tel/Fax:* +33 (0)4.99.61.30.86

Article type Methods

Number of words abstract: 150, main text: 5000

Number of references 61

Number of tables or figures 6

Acknowledgements. We thank Laura Pollock (McGill university) and Tamara MÜNKENMÜLLER (LECA) for guidance on and access to source code for simulating virtual communities; we thank Philippe Choler (LECA) for discussion and crucial explanations on the ecology of Alpine plant communities. We also thank Li Ping Liu for access to source code on exponential family embeddings.

SS is supported by a joint PhD fellowship between the French National Institute of Agricultural and Environmental Research (INRAE) and the French Research Institute for digital sciences (Inria). The research was also supported by the Agence Nationale pour la Recherche (ANR) through the GlobNet (ANR-16-CE02-0009), Gambas (ANR-18-CE02-0025) and Forbic (ANR-18-MPGA-0004) projects. Most of the computations presented in this paper were performed using the GRICAD infrastructure.¹

Authorship. SS conceptualized the inference framework and reviewed it with EG. WT and SS designed the evaluation methodology. SS, WT and EG analyzed the results. MH gave additional perspectives to the paper. SS wrote the first version of the paper. All authors contributed critically to editing the manuscript and gave final approval for publication.

Data Availability. Data on Alpine plants [13], obtained from the *ade4* package. Simulated datasets and source codes are available in the online GitHub repository.

Source code. The source code for running the model is released as a Python package available at DOI:10.5281/zenodo.3611437. A tutorial for running it from R is provided in the Supplementary Materials.

¹<https://gricad.univ-grenoble-alpes.fr>

Abstract

The interplay between environmental suitability, dispersal and biotic interactions induces spatial patterns of species' co-abundance. Existing statistical frameworks that infer the underlying interactions from these patterns either ignore the species response to the environment or they fail to account for the asymmetric nature of interactions.

Here, we propose a framework that (a) models pair-wise associations as directed influences from a source to a target species, parameterized with two species-specific latent variables: the response of the target species to the community, and the effect of the source species on the community; and (b) jointly fits these associations with a habitat suitability model through a conditional abundance model. Using both simulated and empirical data, we demonstrate the ability of the framework to recover known associations and highlight the properties of the learned association networks. Our framework should now pave the way for getting more accurate pictures of interspecific dependencies from empirical data.

1 Introduction

Understanding the drivers of species distribution and their abundance is a long-lasting goal of biogeography [32]. Niche theory explains the spatial distribution of species by a set of physiological and adaptive properties allowing them to thrive in specific environmental conditions and decline in others [12, 49]. The range of environmental (abiotic) variables, such as climate and soil characteristics, that matches the eco-physiological requirements of a species delimits its potential niche (Grinnellian niche, [25]) or (Fundamental niche, [33]). Habitat suitability models (HSM) or species distribution models (SDMs) [27] aim to infer and model this niche by establishing statistical relationships between observed occurrences or abundances of species and the environmental characteristics of the corresponding locations.

HSMs have proven useful to predict species ranges in space, or their shift in response to climate change, providing operational tools to conservation biologists [19,26]. However, as they model multiple species distributions separately, they fail to detect, or to account for, possible dependencies between species that can restrict or extend their ranges beyond what is expected when considering only abiotic factors. Indeed, species may exclude one another locally (*competitive exclusion*, [28]) or be different enough in terms of space and resource needs to co-exist (*niche partitioning*, [52]). Conversely, some species facilitate others by modifying the environment in a way that creates habitats or enables access to resources for other species (*engineering and facilitation*, [17]). Although these interactions take place at a local scale, some of them may alter the range of the species on a wider, macroscopic scale [22]. Consequently, they induce consistent patterns of co-location and dislocation that are unexplained by the abiotic environment. Such interactions are referred here as *associations*. The inability to take into account the presence or absence of other species is therefore an important source of errors for statistical models of species distributions ([61]).

Over the last decade, several approaches have been proposed to infer interspecific dependencies from the observations of many species. Probabilistic Graphical Models (PGM) [34] have been used to infer either directed (Bayesian Networks, BN) or undirected (Markov Random Fields, MRF) networks [20] involving plants [1], parasites and potential hosts [46], predators and preys [37,56] or multi-trophic communities turnover [43]. More recently, Joint Species Distribution Models (JSDM) were introduced to address the same question while jointly predicting co-occurrences of multiple species [44, 48]. The gist is that once abiotic factors are accounted for, the unexplained variance, typically captured by the correlation matrix of the residuals, is attributed to the effect of species on one another or to unknown environmental variables [44]. As they rely purely on correlations, JSDMs and MRFs are limited to estimating symmetric associations where the involved parties influence one another with the same polarity and strength. Lany et al. [35] proposed a JSDM that allows to capture asymmetric associations but requires longitudinal data (see also [2, 30]). On the other hand, BNs support directed relationships but they impose an acyclic structure that does not allow modeling of bidirectional influences.

Inferring associations from co-occurrence data is a common task in text mining. Supposedly, the probability of a word occurring in a particular sentence of a text depends on the semantic compatibility (association) of this word with the list of words surrounding it, forming its context. The common approach is to use *word embedding* algorithms (e.g *word2vec* [38]) to learn multidimensional representations (embeddings) of words that encode this

38 contextual semantic compatibility. By analogy, in community ecology, the probability of the presence of a species
39 in a given environmentally suitable site depends on its compatibility with other species occurring at that site, i.e.
40 other species in the observed community.

41 Recently, word embeddings were generalized to any type of data that follow an exponential family distribution
42 [36, 50], including binary and ordinal data, in so-called exponential family embeddings. Building on this work,
43 we propose here a conditional probabilistic model of species co-distributions that can be trained jointly with any
44 habitat suitability model on presence/absence or count data to infer interspecific associations. In this paper,
45 we detail the methodology, the mathematical formulation and the underlying assumptions. We evaluate the
46 capacity of the model to accurately recover associations using both simulated data for which we know the true
47 interactions between species, and empirical data for which reliable expert knowledge is available [13, 60]. We show
48 that our model critically infers meaningful associations. Finally, we demonstrate how the learnt parameters can
49 be harnessed, to analyze the structure of biotic association networks.

50 2 The inference framework

51 Three main conditions should be satisfied for a species to be present at a given site. First, the site must be
52 accessible. This relates to the species' intrinsic *dispersal* capacity and the presence of migration opportunities
53 or barriers. Second, the abiotic conditions should allow the species' population to maintain a positive growth
54 rate. This condition is referred to here as *habitat suitability* and is the target of Habitat Suitability Models.
55 Third, the species should sustain the interactions with the other species of the community, since those interactions
56 can also impact the species' survival chances and its abundance [27]. Although we recognize the importance of
57 spatial dispersal processes, in this study we focus on the latter two factors, namely *habitat suitability* and *species*
58 *interactions*.

59 **Notation.** We consider a dataset consisting of the abundances of a collection \mathcal{S} of m species observed at a
60 collection \mathcal{K} of n sites, as well as abiotic variables measured at these same sites or in their vicinity. The abundance
61 of species i at site k is denoted y_{ki} , while the vector x_k represents the *abiotic variables* at site k .

62 In what follows, we introduce the key concepts used in the inference model. In particular, we explain how we
63 model the associations between a pair of species by decomposing them into effects and responses, represented as
64 multi-dimensional embedding vectors, and how we use these embeddings to recover biotic interactions.

65 2.1 Spatial associations and biotic context

66 2.1.1 Representing species associations using embeddings

67 For a given pair of species, a *spatial association* describes the relative influence that they have on each other's
68 abundance. The two directions of this influence can be of different types (positive, negative or neutral) and have
69 different intensities Fig 1b. Several mechanisms can lead to such association: a direct interaction between these two
70 species (e.g. competition, predator-prey), an indirect interaction through the environment or a shared correlation

71 to an unmeasured environmental variable or an unobserved group of organisms.

72 Here, we represent the association between species i and j with a pair of scalars a_{ij} and a_{ji} , representing the
73 strength of the influence of species j on species i and vice-versa, respectively. More specifically, a_{ij} represents
74 the change (excess if positive, deficit if negative, none otherwise) in *target* species i 's abundance induced by the
75 *source* species j . These values across all pairs of species, and in both directions, can be collected into an $m \times m$
76 asymmetric association matrix A .

77 The association strength depends on two parameters: the *effect* applied by the source on the target species,
78 and the *response* of the target species. We assume these parameters are controlled by intrinsic traits or properties
79 of the species, which we encode in two separate d -dimensional real-valued vectors referred to as embeddings. In
80 practice, d is a user defined hyperparameter which is typically significantly smaller than half the number of species.

81 The *effect embedding* of species i , α_i , captures the type of organisms the species allows when it is present. The
82 *response embedding* of species i , ρ_i , measures the type of biotic context the species would strive in. For instance,
83 trees with spreading canopy create shade (effect) that selects only shade-tolerant (response) species and exclude
84 others. The response and effect embeddings of the different species can be collected into two $m \times d$ matrices,
85 respectively denoted as P and Q .

86 The association matrix Fig 1a is then written as $A = PQ^T$.

87 2.1.2 Biotic context

88 The biotic context encodes our assumptions about the potential biotic effects a target species is exposed to at a
89 given site. In the simplest case, without any prior knowledge, it consists of individuals from other species observed
90 at the same site. Formally, the biotic context of species i at site k , denoted C_{ki} , is defined as follows:

$$C_{ki} = \{j \in \mathcal{S}, j \neq i \text{ and } y_{kj} > 0\}.$$

91 We obtain the aggregated effect of the biotic context by averaging the effect embeddings of its elements weighted
92 by their respective abundances:

$$z_{ki} = \frac{1}{|C_{ki}|} \sum_{j \in C_{ki}} y_{kj} \alpha_j.$$

93 This formulation allows the presence of facilitators and competitors to balance one another. By weighting with
94 abundance, we implicitly consider that individuals from the same species are similar and contribute equally to the
95 community structure. Rare species have a noticeable impact only if their per capita effect is stronger than the
96 aggregated effect of dominant groups.

97 The biotic context carries implicit constraints on the structure of species association networks by restricting
98 the set of potential associations a priori. For instance, it can be customized for each species according to its
99 known interactions. Moreover, it can include species from neighboring locations (**spatially-explicit**) up to a
100 chosen radius where their influence is relevant (e.g. species with high mobility). Similarly, we can construct
101 the biotic context from previous observations (**temporally-explicit**) to perform a causal analysis. We detail the

102 mathematical adjustments of these alternative definitions along with the associated data requirements and relevant
 103 effect aggregation functions in the Appendix.

104 2.2 A conditional generative model of abundance

105 2.2.1 Formalization

106 The habitat suitability for species i at site k , denoted s_{ki} , is a binary variable that follows a Bernoulli distribution
 107 whose parameter (success rate) is estimated using a habitat suitability model (HSM), h , fitted on the target
 108 species's occurrences, i.e.

$$s_{ki} \sim \mathcal{B}(h_i(x_k)).$$

109 At sites where the abiotic environment is unsuitable (i.e. where $s_{ki} = 0$), the probability mass of the species
 110 abundance is concentrated on zero. Otherwise (i.e. where $s_{ki} = 1$), the abundance of the species is a function of
 111 its biotic context. In other words, we assume that the abiotic environment conditions the presence or the absence
 112 of a given species, while the biotic context controls its abundance.

113 Following Rudolph et al. [50], we model the abundance of species i at site k using the canonical form of the
 114 exponential family \mathcal{E} , whose probability density function (pdf) $f_{\mathcal{E}}$ is parameterized by η_{ki} . Formally,

$$y_{ki} \sim \begin{cases} \mathcal{E}(\eta_{ki}, \tau(y_{ki})) & \text{if } s_{ki} = 1, \\ \delta_0 & \text{otherwise,} \end{cases}$$

115 such that

- 116 - δ_0 denotes Dirac (point-mass) distribution, whose density is equal to one at zero, and to zero elsewhere.
- 117 - τ denotes the sufficient statistic from the canonical form of the exponential family distribution. It depends
 118 only on the data point y_{ki} .

119 We let the canonical parameter η_{ki} depend on the response ρ_i of the target species and on the biotic context
 120 effect z_{ki} . An offset o_i is used to represent the baseline abundance of each species in the event of an empty biotic
 121 context. The link function f scales the outcome to the domain of the target variable. The canonical parameter
 122 η_{ki} is defined as

$$\eta_{ki} = f(\rho_i z_{ki} + o_i),$$

123 which can be rewritten as an aggregate of pairwise association strengths:

$$\eta_{ki} = f\left(\sum_{j \in C_{ki}} y_{kj} a_{ij} + o_i\right).$$

124 The type of data considered (presence/absence vs. abundance) might lead to different choices of probability
 125 distributions, which in turn require resorting to different variants of the generic model (cf. Table 1).

126 2.2.2 Inference

127 We gather as positive and negative observations (instances), the species that are present and absent at each site.
128 Absent species are typically over-represented in an ecological dataset as compared to present species, leading to a
129 high imbalance between positive and negative observations. To address this issue, while all positive instances are
130 included into the training set, negative instances are sub-sampled randomly (at rate $r\%$) at each training iteration.
131 Then, we use Stochastic Gradient Descent [6] to learn the parameters (esp. response and effect embeddings matrices
132 and HSM parameters) that minimize the negative loglikelihood (cf. Appendix) of the selected observations with
133 the addition of lasso penalties on the embeddings to promote the sparsity of the resulting associations. Finally, we
134 use cross-validation to select the hyper-parameters (including hyper-parameters for the abiotic suitability model,
135 embedding dimension, vector of species offsets, negative examples subsample rate and regularization coefficient).

136 2.3 Unraveling inter-specific association networks

137 To identify meaningful associations, we apply two filtering steps to the estimated matrix A . First, the *statistical*
138 *filtering* step consists in setting to zero all associations with a confidence interval containing zero and keeping the
139 mean value for the rest. Second, the *biogeographic filtering* step aims to further eliminate associations that are
140 predicted to potentially exist from the latent representations of the species, but are not realized by the species
141 occurrences because they break some biogeographic constraints:

- 142 1. Mutualism or attraction between two species require co-existence. Thus, we set to zero any inferred positive
143 effect involving two species never observed together in the same site [51], i.e. non-co-occurring.
- 144 2. Repulsive relationships do not require co-occurrence and may even explain the geographic separation. Hence,
145 the involved species do not have to co-occur but should live in similar environments to be considered as a
146 potential negative association. Specifically, we compute the ranges of the environmental values corresponding
147 to the occurrences of each species and retain negative associations only if these ranges overlap or if the
148 species are otherwise sufficiently similar (above a user-defined similarity threshold) in terms of their habitat
149 suitability parameters.

150 Furthermore, we focus on the polarity of the associations, rather than their strength, hence we consider a
151 discrete version of the association matrix, which we call the *adjacency matrix* defined as follows:

$$I_{ij} = \begin{cases} \text{positive} & \text{if } a_{ij} > \epsilon^+, \\ \text{negative} & \text{if } a_{ij} < \epsilon^-, \\ \text{neutral} & \text{otherwise.} \end{cases}$$

152 such that ϵ^+ and ϵ^- represent a user-defined threshold on the strength of the positive and negative associations,
153 respectively.

154 The resulting matrix defines a network, where each species is represented by a vertex and a directed edge
155 labelled as positive (resp. negative) from vertex i to vertex j represents a positive (resp. negative) influence of
156 species i on species j .

157 By design (Fig. 1a), species with similar response embeddings constitute clusters of rows in the adjacency
158 matrix, called *response groups*. Conversely, species with similar effect embeddings constitute clusters of columns
159 called *effect groups*. These groups can be computed simultaneously using a *bi-clustering* algorithm [23]. The
160 product of both types of groups results in the emergence of clusters of exchangeable or redundant species in the
161 resulting network, called *structural roles* [21].

162 **3 Theoretical validation of the framework**

163 Before applying our model to infer associations from empirical data, we evaluated its ability to recover interspecific
164 associations from simulated datasets with known associations.

165 **3.1 Data generation**

166 We used a process-based stochastic model adapted from `Virtualcom` [40] to simulate the assembly of individuals
167 from a regional species pool into communities, on different locations sampled along an environmental gradient.
168 The assembly process is controlled by three filtering mechanisms: the response to the abiotic environment, the
169 outcome of biotic interactions, and reproduction.

170 We set up an experiment where multiple simulations were run on random points on a single abiotic gradient
171 ranging from 0 to 100 with different hand-crafted configurations of the prior association matrix: absence of associ-
172 ation, positive associations only, negative associations only and both positive and negative associations. Generally,
173 few species were set to interact. We investigated two settings where species are involved in one association (sparse)
174 or more than one associations (dense). Associated pairs were chosen such that their abiotic niches overlap.

175 In each configuration mode, we varied the number of species (5, 10 or 20), the density setting (sparse or dense)
176 and whether the association matrix included asymmetric effects (semi-attraction or semi-repulsion). Positive (resp.
177 negative) effects were all set to +1 (resp. -1) as we are interested in the polarity of the associations rather than
178 their magnitude. The factorial design of this experiment produced 29 simulation datasets (Appendix).

179 **3.2 Evaluation**

180 We describe the inference procedure and model selection in detail in Appendix.

181 **3.2.1 Observed relative abundances vs. inferred associations strengths**

182 For each association type and simulation configuration, we compare the inferred association strengths to the
183 observed relative abundance effects. To quantify these effects, we define the *relative abundance index* (RAI), an
184 asymmetric metric that measures the change in abundance of the target species when the source species is present
185 as compared to its mean abundance irrespective of whether the source species is present.

186 Formally, we define

$$\bar{y}_t = \text{avg}(\{y_{kt}, \text{ for all } k \in \mathcal{K} \text{ such that } y_{kt} > 0\}), \quad \text{and}$$

$$\Delta_{st} = \{y_{kt} - \bar{y}_t, \text{ for all } k \in \mathcal{K} \text{ such that } y_{kt} > 0 \text{ and } y_{ks} > 0\}.$$

187 Then $\text{RAI}_{st} = \text{avg}(\Delta_{st})$. The larger the standard deviation $\text{std}(\Delta_{st})$, the more ambiguous the strength of the
 188 effect of species s on species t . If the confidence interval $\text{avg}(\Delta_{st}) \pm 1.96 \text{std}(\Delta_{st})$ does not contain zero, then
 189 the simulated dependencies unambiguously translate a polarized effect of species s on species t . Otherwise, the
 190 polarity of the effect is ambiguous, due to either confounding effects of other species or a neutral association if
 191 the mean is close to zero. We also compute the Jaccard coefficient between the binary presence/absence vectors
 192 of species s and t , a.k.a. Jaccard co-occurrence index, denoted as J_{st} .

193 3.2.2 Association classification evaluation

194 We discretized the learnt associations using the threshold values $(\epsilon^+, \epsilon^-) = (0, 0)$ to obtain the corresponding
 195 classes (positive, negative, neutral). Subsequently, we evaluated the latter against the simulated association classes
 196 as the ground-truth using standard multi-class performance metrics (recall, precision, F1-score). Recall measures
 197 the percentage of associations of a specific class correctly recovered by the model, whereas precision quantifies
 198 the percentage of true associations amongst those classified as the specified class. The F1-score is computed as
 199 the harmonic mean of recall and precision. A higher recall indicates lower false negatives whilst higher precision
 200 indicates lower false positives.

201 3.3 Results

202 Overall, we found a better fit for positively associated communities than those in competition, while mixing
 203 both types resulted in intermediate performances. In the case of simulations with competition, most sparse
 204 and asymmetric configurations induced better performances than their (dense and symmetric) counterparts (see
 205 Appendix).

206 The average relative abundance index reflected well the simulated associations with negative (resp. positive)
 207 effects below or around (resp. above) zero, while neutral associations were centered around zero. However, most
 208 positive effects yielded small relative abundance effects as compared to the negative effects. Although more clearly
 209 marked, the latter approached neutrality on larger and more densely connected communities (Fig. 2).

210 The inference model was able to discriminate positive from negative effects while maintaining an average value
 211 for non interacting pairs centered on zero with a small variance. On simulations with a dense mix of positive and
 212 negative associations, both observed effects and inferred associations were close to zero, possibly due to opposing
 213 effects canceling each other. The absence of associations led to the systematic prediction of the offset hence the
 214 constant deviance on mixed types simulations.

215 On average, recall did not vary significantly between positive and negative associations, whereas precision
 216 was higher for negative than for positive associations (Table 2), indicating the detection of spurious positive

217 associations. Much higher precision was achieved for neutral associations (absence of association). The prediction
218 performance was better for small species pool, with higher recall on the dense configurations but higher precision
219 on sparse ones. The sparse asymmetric positive simulation resulted in the worst predictive performance with low
220 recall.

221 4 Empirical case study

222 We applied our model to empirical observations (also analyzed by [60]) of species abundances to show how it can
223 unravel meaningful associations. We used the plant dataset from [13] that consists of 75 vegetation plots, of size
224 5×5 m sampled around July 2000. Across the vegetation plots, the abundance of the 82 occurring plant species
225 were registered. In addition, a set of environmental and topographic variables was recorded on each plot. We
226 describe the data preprocessing, framework adaptation and model selection procedures in Appendix.

227 4.1 Network analysis

228 We performed a hierarchical bi-clustering on the inferred association matrix to obtain effect and response groups
229 (Fig. 3a). In parallel, we applied the modularity maximization algorithm [42] on the association network to identify
230 densely connected modules, referred to as communities [21]. After that, we mapped the structural roles within
231 the modules to create the group-level network. Finally, we analyzed the resulting patterns in light of existing
232 literature on alpine plants interactions [14].

233 4.2 Results

234 We identified four densely connected modules of different sizes, within which species occupied various structural
235 roles in the plant association network. Modules were structured along the snow melting date gradient (Fig. 3a).

236 Species from early-melting sites were classified into the same module. We found a prominence of positive asso-
237 ciations, specifically an unselective mostly asymmetric attraction of forbs and grasses to tall dominant graminoids
238 (*Carex*, *Kobresia*). Forbs and grasses formed two distinct groups linked by negative associations. Besides, some
239 of them acted as hubs connecting the high elevation sites to the adjacent sites where they also occurred. The
240 second module encompassed two groups of grasses: (i) Tall herb grasslands occurring in favorables conditions,
241 mostly structured by negative associations (ammensalism and competition); (ii) Short herb meadows, exposed to
242 zoogenic disturbances. They presented higher abundances when co-occurring with tall herbs. The third module
243 consisted of chinophilous (cold-resistant) vegetation appearing on late-melting sites. The last module included
244 north-facing isolated communities dominated by *Salix Herbacea* positively associated with high-altitude commu-
245 nities and characterized by high eccentricity (Fig. 4).

246 In general, positive associations were prominent on stressful conditions. For instance, on early melting sites,
247 species are exposed to wind and erosion due to snow melting [13]. The positive associations could be explained
248 by the facilitative effect of graminoids through multiple hypothetical mechanisms. Graminoids have the ability to
249 maintain the soil stability [9,29], they can also prevent dessication and frost heaving on stones in favor of seedling

250 survival [14]. They also sustain a suitable microclimate for small forbs and grasses, while offering protection from
251 the wind [59]. On the other hand, negative associations were found on the richest sites, hypothetically reflecting
252 a competition for resources: water and Nitrogen [14].

253 As reported in the literature, the abiotic conditions strongly structured the predicted response [13] of the plant
254 species and the dominant interaction types [10]. Specifically, network modules were distributed along the gradient
255 following their composition’s response to the average snow duration. Negative associations inflicted by competitive
256 tall grasses on mid slope communities connected early-melting communities to the chinophilous vegetation from
257 late-melting sites in the resulting association network (Fig. 3b). Response and effect groups were included in one or
258 at most two close (in terms of position in the gradient) modules (Fig. 3a). Non neutral associations had consistent
259 types (either negative or positive) within effect groups regardless of the responding species, suggesting that species
260 roles (effects) in their community might be predictable from their own characteristics and the surrounding abiotic
261 conditions. At last, learnt associations were symmetric within groups but asymmetric (mostly semi-attraction or
262 semi-repulsion) between them.

263 5 Discussion

264 In this article, we tackled the challenge of inferring interspecific associations from multiple species co-abundances
265 along environmental gradients. To do so, we formalized pairwise associations as a function of two sets of latent
266 variables representing the response and the effect of each species in respect to the others. We incorporated these
267 associations into a conditional probabilistic model of abundance that controls for habitat suitability. The evaluation
268 of the model’s ability to recover known associations from simulated data showed that it is able to discriminate the
269 association types (positive, negative or neutral), but the inferred strength depended on the species pool size, niche
270 overlap, network density and the presence of multiple confounding associations. When we applied the model to
271 the co-abundance data of plants along a mesotopographic gradient in the French Alps, the model identified most
272 of the important relationships expected in these plant communities [14].

273 5.1 Disentangling confounding effects

274 Inferring species associations from co-occurrence patterns is a very challenging task [11]. Even in our simulated
275 dataset for which we have pre-defined interactions between species, the resulting co-occurrence levels could be high
276 even for known competing pairs of species (case of a small species pool with large carrying capacity). Instead, the
277 co-abundance structure better reflects the nature of associations. Indeed, species’ abundances are lower than in
278 average in presence of negative associations and higher with positive associates. Nonetheless, pairwise abundance
279 effects may turn out to be neutral in presence of multiple confounding effects [7]. Our proposed framework is able
280 to tease apart the different opposing influences (Fig. 2) by estimating the pairwise associations conditioned on all
281 other species.

5.2 Interaction of habitat suitability and biotic associations

Following a hierarchical filtering scheme of community assembly [7, 44, 45], we assumed the species occurred only on suitable habitats. However, species may occur outside their abiotic niche if the local conditions are ameliorated by their peers (facilitation) [8, 16, 53].

In the absence of micro habitat descriptors, two scenarios are possible depending on the location’s distance in environmental space from the species’ true fundamental niche. If the distance is small, any HSM would probably overestimate the species ranges [5] with no residuals on the occurrence probabilities. Consequently, the associations would go undetected by joint presence/absence models (e.g. JSDMs), whereas the proposed model will detect them as long as there is an observed increase in abundance. Conversely, if the distance is large then a robust HSM would correctly learn the true abiotic niche, leaving the unexplained presence to the biotic effect. But since the habitat was found unsuitable by the HSM component, our model would fail to detect the underlying biotic association.

Indeed, the model constraints considered co-occurrence as a prerequisite for positive associations (including facilitation) while in the described scenario co-occurrence is a consequence of facilitation. A compromise solution would be to consider that species do not respond separately to each factor : abiotic captured by the HSM parameters θ_h and biotic through the response embeddings ρ . Instead, they respond to the outcome of the aggregation f of the abiotic features x and the biotic context effects z . By choosing f to be a universal function approximator (e.g a feed-forward neural network [31]), we can learn which of the hierarchical, additive or interactive models of the abiotic and biotic filters best fit the observed abundances/distributions without any *a priori* assumption.

5.3 Associations validation

With no prior or guidance on the expected associations network, validating inferred associations is challenging.

It is now agreed upon that associations are not equivalent to biological interactions [18]. They represent significant spatial co-location patterns, that are informative in a predictive rather than in a causal way [39]. The specific mechanism that led to these patterns may vary from pair to pair, ranging from direct interactions (e.g. trophic), to indirect interactions (e.g. engineering, shared habitat). Consequently, the validation of inferred associations should consider other explanations than biotic interactions.

In general, because many processes influence community assembly, multiple scenarios could lead to the same communities making this problem unidentifiable [41]. In this case, rather than a single expected list of associations, we need all the possible combinations of associations, or a goodness-of-fit measure that accounts for equivalence between different combinations. A possible way to prevent this issue is to include prior knowledge of ecological interactions in the model [11]. For instance, [54] used a Bayesian network with a predefined structure, and trained its parameters using a HSM to predict species occurrence probabilities. In our case, such constraints can be defined by altering the biotic context definition. One direct way to do it is to consider a customized biotic context for each species composed of the set of its potential interaction partners in a pre-built regional metaweb.

There is now growing evidence that ecological interactions are context-dependent [47, 55], we showed in the Appendix how to adapt the framework to infer associations whose strength is modulated by other covariates (e.g. stress, presence of predator, etc.). Recently developed models account for association variability as a function of

318 the environmental context [15, 55]. Despite these new possibilities, the question of how to validate their results
319 still arises itself.

320 5.4 Species roles and association’s asymmetry

321 In the plant network, modules were strongly structured by the abiotic gradient. Because of the HSM conditioning,
322 this pattern would be expected in gradients with strong taxonomic turnover. Modules and structural roles provided
323 two complementary information. The modules brought insight into the connectivity therefore the stability [24]
324 of the meta-network while structural roles, including response and/or effect groups, were useful to evaluate the
325 functional redundancy within locally projected networks. In the future, identifying characteristic traits within
326 structural roles would allow the elicitation of the functional drivers of network structure.

327 Many studies [57] of interaction networks reported that interactions tend to be asymmetric, both in terms of
328 “type” in binary networks [3, 58] and in terms of “strength” in quantitative networks [4]. By analyzing various
329 types of observed ecological networks, [57] suggested that asymmetric interaction strengths arise from mismatch in
330 species relative abundances. Since the proposed model learns per capita effects, the abundance is already controlled
331 for. In the case of Alpine plants, the response and effect groups encompassed species occurring in similar habitats.
332 Knowing that, the predominance of symmetrical (resp. asymmetric) associations within (resp. between) groups
333 suggests the degree of asymmetry might be inversely related to habitat sharing or niche overlap.

334 Nevertheless, the ability to discern this asymmetry sheds light on the imbalance and direction of interspecific
335 dependencies, drawing a more accurate picture for biodiversity forecasting models.

336 6 Conclusion

337 Biological interactions and other processes induce spatial patterns of co-occurrence and co-abundance. We pre-
338 sented and validated a model of species co-abundances as a function of the habitat and biotic associations. We
339 proposed an asymmetric scheme for modeling associations that is based on learning latent representations of
340 species’ responses and effects. Future efforts should be directed towards a combination of prior knowledge on the
341 complete or partial topology of the association networks to guide the inference process. Along with that, a strong
342 theory of how known ecological interactions influence the co-distribution of species is needed to support all these
343 models.

344 References

- 345 [1] ADERHOLD, A., HUSMEIER, D., LENNON, J. J., BEALE, C. M., AND SMITH, V. A. Hierarchical bayesian
346 models in ecology: reconstructing species interaction networks from non-homogeneous species abundance
347 data. *Ecological Informatics* 11 (2012), 55–64.
- 348 [2] BARRAQUAND, F., PICOCHÉ, C., DETTO, M., AND HARTIG, F. Inferring species interactions using granger
349 causality and convergent cross mapping. *arXiv preprint arXiv:1909.00731* (2019).
- 350 [3] BASCOMPTE, J., JORDANO, P., MELIÁN, C. J., AND OLESEN, J. M. The nested assembly of plant–animal
351 mutualistic networks. *Proceedings of the National Academy of Sciences* 100, 16 (2003), 9383–9387.
- 352 [4] BASCOMPTE, J., JORDANO, P., AND OLESEN, J. M. Asymmetric coevolutionary networks facilitate biodi-
353 versity maintenance. *Science* 312, 5772 (2006), 431–433.
- 354 [5] BEALE, C. M., AND LENNON, J. J. Incorporating uncertainty in predictive species distribution modelling.
355 *Philosophical Transactions of the Royal Society B: Biological Sciences* 367, 1586 (2012), 247–258.
- 356 [6] BOTTOU, L. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMP-*
357 *STAT’2010*. Springer, 2010, pp. 177–186.
- 358 [7] BOULANGEAT, I., GRAVEL, D., AND THULLER, W. Accounting for dispersal and biotic interactions to
359 disentangle the drivers of species distributions and their abundances. *Ecology letters* 15, 6 (2012), 584–593.
- 360 [8] BULLERI, F., BRUNO, J. F., SILLIMAN, B. R., AND STACHOWICZ, J. J. Facilitation and the niche: impli-
361 cations for coexistence, range shifts and ecosystem functioning. *Functional Ecology* 30, 1 (2016), 70–78.
- 362 [9] CALLAWAY, R. M. *Positive interactions and interdependence in plant communities*. Springer, 2007.
- 363 [10] CALLAWAY, R. M., BROOKER, R., CHOLER, P., KIKVIDZE, Z., LORTIE, C. J., MICHALET, R., PAOLINI,
364 L., PUGNAIRE, F. I., NEWINGHAM, B., ASCHEHOUG, E. T., ET AL. Positive interactions among alpine
365 plants increase with stress. *Nature* 417, 6891 (2002), 844.
- 366 [11] CAZELLES, K., ARAÚJO, M. B., MOUQUET, N., AND GRAVEL, D. A theory for species co-occurrence in
367 interaction networks. *Theoretical Ecology* 9, 1 (2016), 39–48.
- 368 [12] CHASE, J. M., AND LEIBOLD, M. A. *Ecological niches: linking classical and contemporary approaches*.
369 University of Chicago Press, 2003.
- 370 [13] CHOLER, P. Consistent shifts in alpine plant traits along a mesotopographical gradient. *Arctic, Antarctic,*
371 *and Alpine Research* 37, 4 (2005), 444–453.
- 372 [14] CHOLER, P., MICHALET, R., AND CALLAWAY, R. M. Facilitation and competition on gradients in alpine
373 plant communities. *Ecology* 82, 12 (2001), 3295–3308.
- 374 [15] CLARK, N. J., WELLS, K., AND LINDBERG, O. Unravelling changing interspecific interactions across envi-
375 ronmental gradients using markov random fields. *Ecology* 99, 6 (2018), 1277–1283.

- 376 [16] CROTTY, S. M., AND BERTNESS, M. D. Positive interactions expand habitat use and the realized niches of
377 sympatric species. *Ecology* 96, 10 (2015), 2575–2582.
- 378 [17] CUDDINGTON, K., BYERS, J. E., WILSON, W. G., AND HASTINGS, A. *Ecosystem engineers: plants to*
379 *protists*, vol. 4. Academic Press, 2011.
- 380 [18] DORMANN, C. F., BOBROWSKI, M., DEHLING, D. M., HARRIS, D. J., HARTIG, F., LISCHKE, H.,
381 MORETTI, M. D., PAGEL, J., PINKERT, S., SCHLEUNING, M., ET AL. Biotic interactions in species dis-
382 tribution modelling: 10 questions to guide interpretation and avoid false conclusions. *Global ecology and*
383 *biogeography* 27, 9 (2018), 1004–1016.
- 384 [19] ELITH, J., AND LEATHWICK, J. R. Species distribution models: ecological explanation and prediction across
385 space and time. *Annual review of ecology, evolution, and systematics* 40 (2009), 677–697.
- 386 [20] FAISAL, A., DONDELINGER, F., HUSMEIER, D., AND BEALE, C. M. Inferring species interaction networks
387 from species abundance data: A comparative evaluation of various statistical and machine learning methods.
388 *Ecological Informatics* 5, 6 (2010), 451–464.
- 389 [21] GAUZENS, B., THÉBAULT, E., LACROIX, G., AND LEGENDRE, S. Trophic groups and modules: two levels
390 of group detection in food webs. *Journal of The Royal Society Interface* 12, 106 (2015), 20141176.
- 391 [22] GOTELLI, N. J. Ecology: Biodiversity in the scales. *Nature* 419, 6907 (2002), 575.
- 392 [23] GOVAERT, G., AND NADIF, M. *Co-clustering: models, algorithms and applications*. John Wiley & Sons,
393 2013.
- 394 [24] GRILLI, J., ROGERS, T., AND ALLESINA, S. Modularity and stability in ecological communities. *Nature*
395 *communications* 7 (2016), 12031.
- 396 [25] GRINNELL, J. The niche-relationships of the california thrasher. *Auk* 34, 4 (1917), 427–433.
- 397 [26] GUISAN, A., AND THUILLER, W. Predicting species distribution: offering more than simple habitat models.
398 *Ecology letters* 8, 9 (2005), 993–1009.
- 399 [27] GUISAN, A., THUILLER, W., AND ZIMMERMANN, N. E. *Habitat suitability and distribution models: with*
400 *applications in R*. Cambridge University Press, 2017.
- 401 [28] HARDIN, G. The competitive exclusion principle. *science* 131, 3409 (1960), 1292–1297.
- 402 [29] HEILBRONN, T., AND WALTON, D. W. Plant colonization of actively sorted stone stripes in the subantarctic.
403 *Arctic and Alpine Research* 16, 2 (1984), 161–172.
- 404 [30] HOLLINGDALE, E., PÉREZ-BARBERÍA, F. J., AND WALKER, D. M. Inferring symmetric and asymmetric
405 interactions between animals and groups from positional data. *PloS one* 13, 12 (2018).
- 406 [31] HORNIK, K., STINCHCOMBE, M., AND WHITE, H. Multilayer feedforward networks are universal approxi-
407 mators. *Neural networks* 2, 5 (1989), 359–366.

- 408 [32] HUMBOLDT, A. V., BONPLAND, A., ET AL. *Essai sur la géographie des plantes*. Chez Levrault, Schoell et
409 compagnie, libraires, 1805.
- 410 [33] HUTCHINSON, G. The multivariate niche. In *Cold Spring Harbor Symposia on Quantitative Biology* (1957),
411 vol. 22, pp. 415–421.
- 412 [34] KOLLER, D., AND FRIEDMAN, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- 413 [35] LANY, N. K., ZARNETSKE, P. L., SCHLIEP, E. M., SCHAEFFER, R. N., ORIANI, C. M., ORWIG, D. A.,
414 AND PREISSER, E. L. Asymmetric biotic interactions and abiotic niche differences revealed by a dynamic
415 joint species distribution model. *Ecology* 99, 5 (2018), 1018–1023.
- 416 [36] LIU, L.-P., AND BLEI, D. M. Zero-inflated exponential family embeddings. In *Proceedings of the 34th*
417 *International Conference on Machine Learning-Volume 70* (2017), JMLR. org, pp. 2140–2148.
- 418 [37] MAJDI, N., HETTE-TRONQUART, N., AUCLAIR, E., BEC, A., CHOUVELON, T., COGNIE, B., DANGER,
419 M., DECOTTIGNIES, P., DESSIER, A., DESVILETTES, C., ET AL. There’s no harm in having too much: A
420 comprehensive toolbox of methods in trophic ecology. *Food webs* 17 (2018), e00100.
- 421 [38] MIKOLOV, T., SUTSKEVER, I., CHEN, K., CORRADO, G. S., AND DEAN, J. Distributed representations of
422 words and phrases and their compositionality. In *Advances in neural information processing systems* (2013),
423 pp. 3111–3119.
- 424 [39] MILNS, I., BEALE, C. M., AND SMITH, V. A. Revealing ecological networks using bayesian network inference
425 algorithms. *Ecology* 91, 7 (2010), 1892–1899.
- 426 [40] MÜNDEMÜLLER, T., AND GALLIEN, L. Virtualcom: a simulation model for eco-evolutionary community
427 assembly and invasion. *Methods in Ecology and Evolution* 6, 6 (2015), 735–743.
- 428 [41] MUÑOZ-TAMAYO, R., PUILLET, L., DANIEL, J.-B., SAUVANT, D., MARTIN, O., TAGHIPOOR, M., AND
429 BLAVY, P. To be or not to be an identifiable model. is this a relevant question in animal science modelling?
430 *animal* 12, 4 (2018), 701–712.
- 431 [42] NEWMAN, M. E. Modularity and community structure in networks. *Proceedings of the national academy of*
432 *sciences* 103, 23 (2006), 8577–8582.
- 433 [43] OHLMANN, M., MAZEL, F., CHALMANDRIER, L., BEC, S., COISSAC, E., GIELLY, L., PANSU, J.,
434 SCHILLING, V., TABERLET, P., ZINGER, L., ET AL. Mapping the imprint of biotic interactions on β -
435 diversity. *Ecology letters* 21, 11 (2018), 1660–1669.
- 436 [44] OVASKAINEN, O., TIKHONOV, G., NORBERG, A., GUILLAUME BLANCHET, F., DUAN, L., DUNSON, D.,
437 ROSLIN, T., AND ABREGO, N. How to make more out of community data? a conceptual framework and its
438 implementation as models and software. *Ecology Letters* 20, 5 (2017), 561–576.

- 439 [45] PETERSON, A. T., SOBERÓN, J., PEARSON, R. G., ANDERSON, R. P., MARTÍNEZ-MEYER, E., NAKA-
440 MURA, M., AND ARAÚJO, M. B. *Ecological niches and geographic distributions (MPB-49)*, vol. 56. Princeton
441 University Press, 2011.
- 442 [46] PILOSOFF, S., FORTUNA, M. A., COSSON, J.-F., GALAN, M., KITTIPOONG, C., RIBAS, A., SEGAL, E.,
443 KRASNOV, B. R., MORAND, S., AND BASCOMPTE, J. Host–parasite network structure is associated with
444 community-level immunogenetic diversity. *Nature communications* 5 (2014), 5172.
- 445 [47] POISOT, T., STOUFFER, D. B., AND GRAVEL, D. Beyond species: why ecological interaction networks vary
446 through space and time. *Oikos* 124, 3 (2015), 243–251.
- 447 [48] POLLOCK, L. J., TINGLEY, R., MORRIS, W. K., GOLDING, N., O’HARA, R. B., PARRIS, K. M., VESK,
448 P. A., AND MCCARTHY, M. A. Understanding co-occurrence by modelling species simultaneously with a
449 joint species distribution model (jsdm). *Methods in Ecology and Evolution* 5, 5 (2014), 397–406.
- 450 [49] PULLIAM, H. R. On the relationship between niche and distribution. *Ecology letters* 3, 4 (2000), 349–361.
- 451 [50] RUDOLPH, M., RUIZ, F., MANDT, S., AND BLEI, D. Exponential family embeddings. In *Advances in Neural*
452 *Information Processing Systems* (2016), pp. 478–486.
- 453 [51] SANDERSON, J. G., AND PIMM, S. L. *Patterns in Nature: The analysis of species co-occurrences*. University
454 of Chicago Press, 2015.
- 455 [52] SCHOENER, T. W. Resource partitioning in ecological communities. *Science* 185, 4145 (1974), 27–39.
- 456 [53] STACHOWICZ, J. Niche expansion by positive interactions: realizing the fundamentals. a comment on
457 rodriguez-cabal et al. *Ideas in Ecology and Evolution* 5 (2012).
- 458 [54] STANICZENKO, P. P., SIVASUBRAMANIAM, P., SUTTLE, K. B., AND PEARSON, R. G. Linking macroecology
459 and community ecology: refining predictions of species distributions using biotic interaction networks. *Ecology*
460 *letters* 20, 6 (2017), 693–707.
- 461 [55] TIKHONOV, G., ABREGO, N., DUNSON, D., AND OVASKAINEN, O. Using joint species distribution models
462 for evaluating how species-to-species associations depend on the environmental context. *Methods in Ecology*
463 *and Evolution* 8, 4 (2017), 443–452.
- 464 [56] TRIFONOVA, N., KENNY, A., MAXWELL, D., DUPLISEA, D., FERNANDES, J., AND TUCKER, A. Spatio-
465 temporal bayesian network models with latent variables for revealing trophic dynamics and functional networks
466 in fisheries ecology. *Ecological informatics* 30 (2015), 142–158.
- 467 [57] VÁZQUEZ, D. P., MELIÁN, C. J., WILLIAMS, N. M., BLÜTHGEN, N., KRASNOV, B. R., AND POULIN,
468 R. Species abundance and asymmetric interaction strength in ecological networks. *Oikos* 116, 7 (2007),
469 1120–1127.

- 470 [58] VÁZQUEZ, D. P., POULIN, R., KRASNOV, B. R., AND SHENBROT, G. I. Species abundance and the
471 distribution of specialization in host–parasite interaction networks. *Journal of Animal Ecology* 74, 5 (2005),
472 946–955.
- 473 [59] WARDLE, D., BARKER, G., BONNER, K., AND NICHOLSON, K. Can comparative approaches based on
474 plant ecophysiological traits predict the nature of biotic interactions and individual plant species effects in
475 ecosystems? *Journal of ecology* 86, 3 (1998), 405–420.
- 476 [60] WARTON, D. I., BLANCHET, F. G., O’HARA, R. B., OVASKAINEN, O., TASKINEN, S., WALKER, S. C.,
477 AND HUI, F. K. So many variables: joint modeling in community ecology. *Trends in Ecology & Evolution*
478 30, 12 (2015), 766–779.
- 479 [61] WISZ, M. S., POTTIER, J., KISSLING, W. D., PELLISSIER, L., LENOIR, J., DAMGAARD, C. F., DORMANN,
480 C. F., FORCHHAMMER, M. C., GRYTNES, J.-A., GUISAN, A., ET AL. The role of biotic interactions in
481 shaping distributions and realised assemblages of species: implications for species distribution modelling.
482 *Biological reviews* 88, 1 (2013), 15–30.

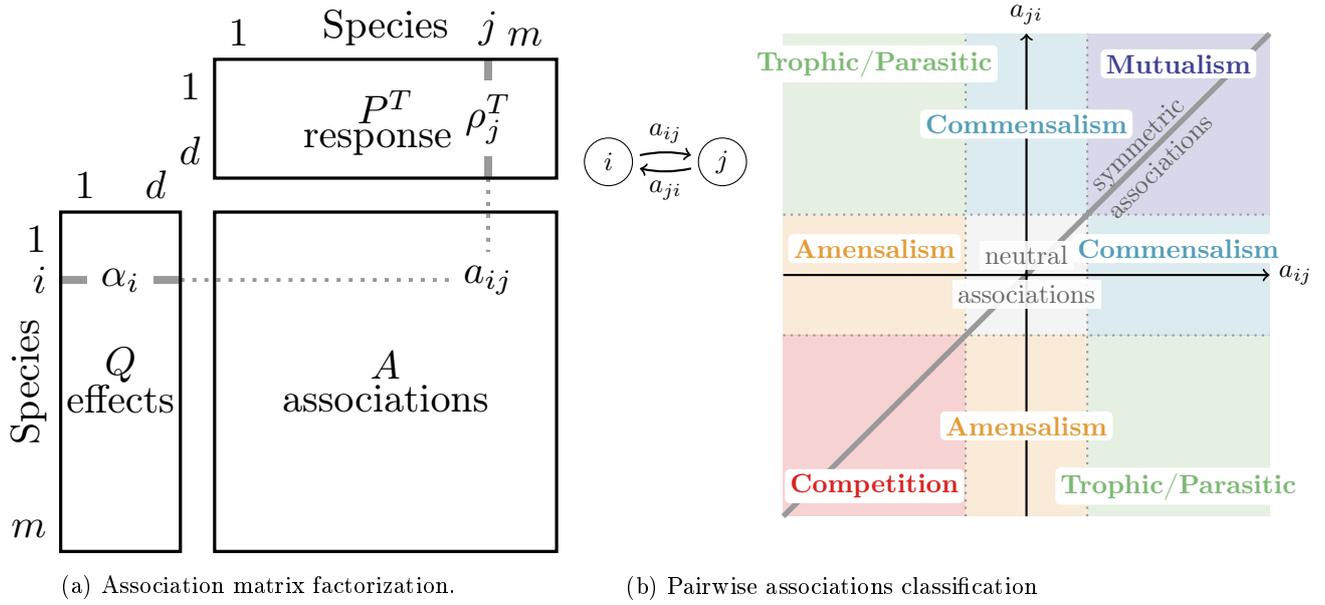


Fig. 1: Association strengths are computed from species response and effects (a). Pairwise association strengths are mapped to potential interaction classes (b). The different quarters of the bi-plot represent the various types of associations between species. The first bisector represents the association domain covered by correlation-based approaches (JSDM) and undirected graphical models (MRF).

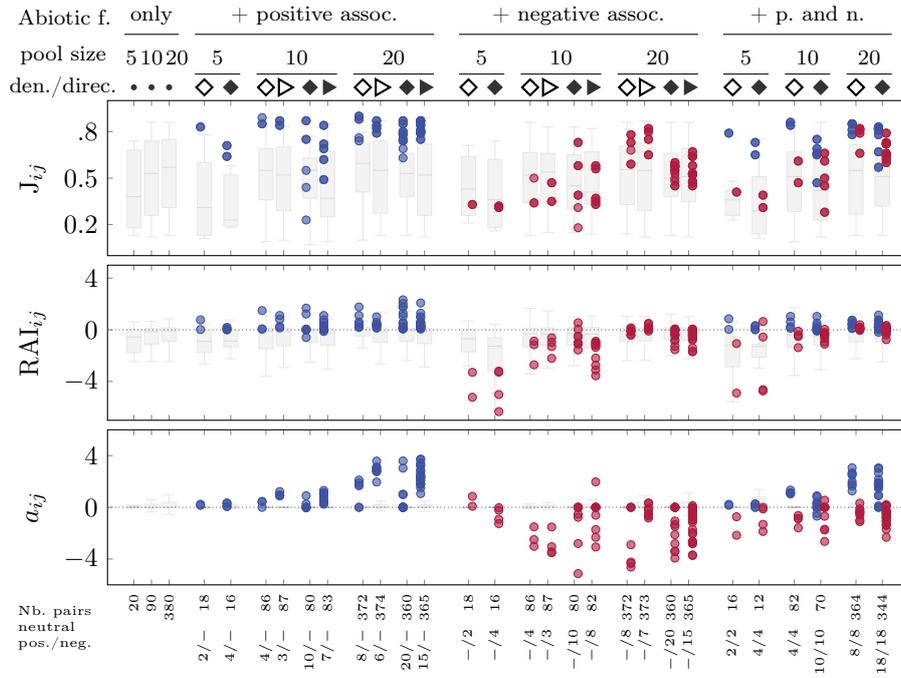
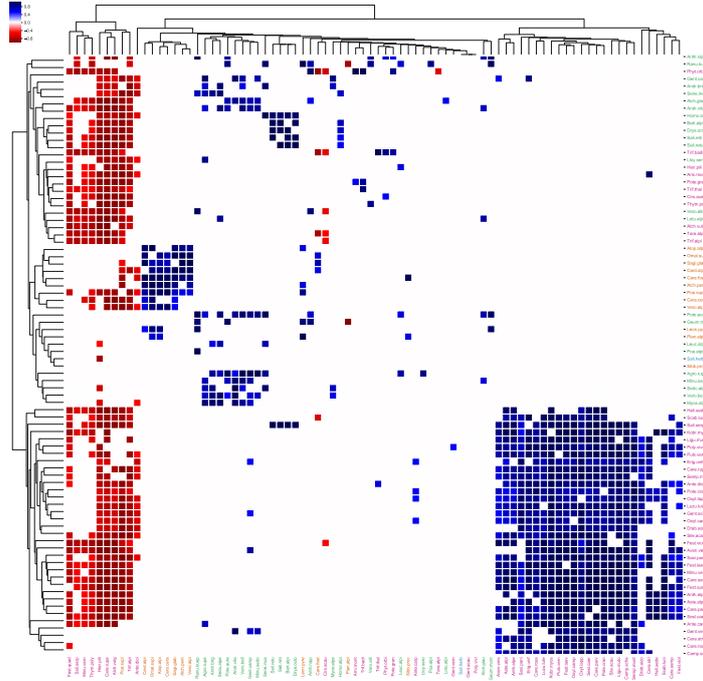
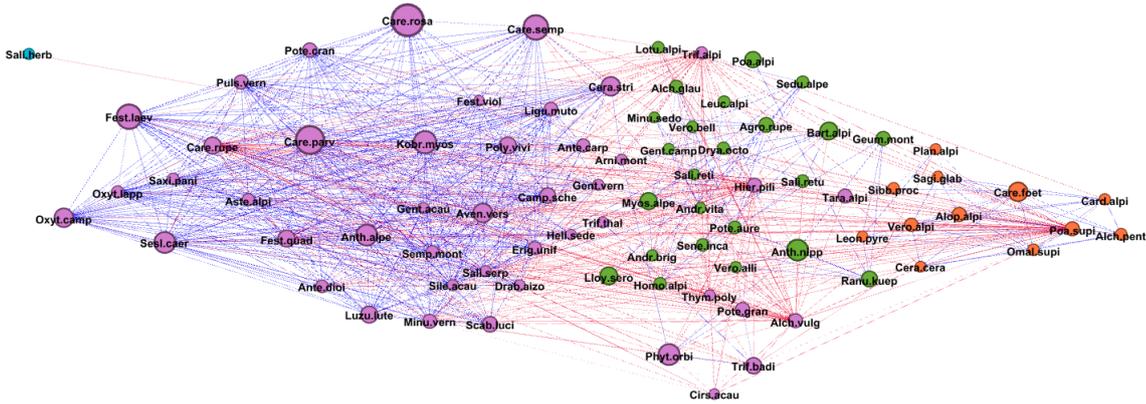


Fig. 2: Distribution of Jaccard co-occurrence indices J_{ij} , relative abundance indices RAI_{ij} , and inferred association strengths a_{ij} per simulation. Each data point represents a directed association (positive in red, negative in blue and neutral in gray) involving two species from the corresponding simulation. For clarity, we represent neutral associations with boxplots. The symbols represent combinations of density (sparse:hollow shape/dense:filled shape) and directionality (asymmetric:triangle/symmetric:diamond)



(a) Inferred plant association matrix. Species in the association matrix are grouped based on a hierarchical bi-clustering performed row-wise (yielding response groups) and column-wise (yielding effect groups).



(b) Network of plant associations. Blue (resp. red) edges indicate negative (resp. positive) edge weights. Node colors on the graph represent communities identified by the modularity maximization algorithm [42] whilst node sizes are scaled according to the plant height. Nodes (except *Salix Herbacea*, which represents the vegetation on the northern face of the gradient) are placed from left to right following an ascending order of their response to Snow duration (regression coefficient from the Generalized Linear Model used as a Habitat Suitability Model).

Fig. 3: Plant associations on an Alpine mesotopographic gradient. We highlight the communities (node colors) in figure (b) using colored labels on the matrix (a).

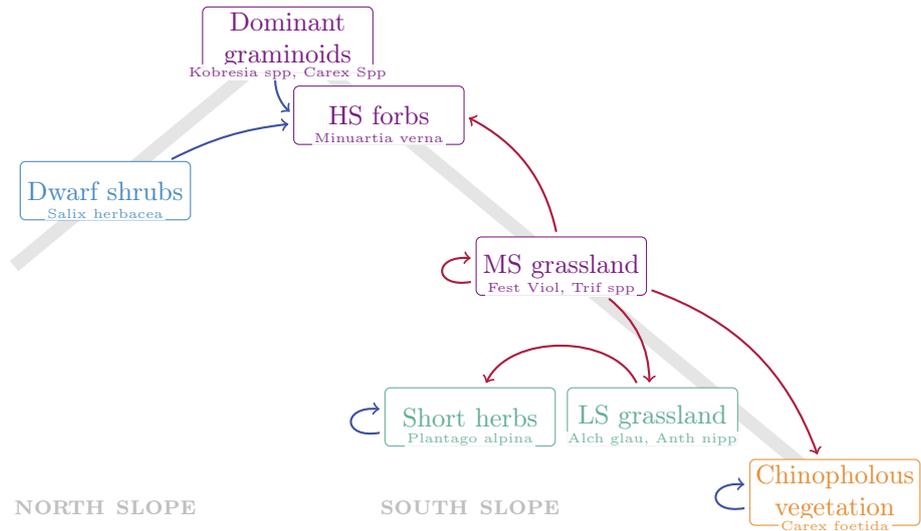


Fig. 4: The summary association network. Structural roles (nodes) are mapped to position in the gradient (Higher-slope HS, mid-slope MS, lower-slope LS) and plant classes (graminoids, grasses/herbs, forbs) and network modules (node colors). Edges go from a source (effect group) to a target (response group). Blue (resp. red) edges represent positive (resp. negative) associations.

Data type	Distribution	Link function	Natural parameter mapping
Presence/ Absence	Binomial	identity	Probability of occurrence $p_{kj} = \sigma(\sum_{i \in C_{ki}} y_{kj} a_{ij} + o_i)$ σ : logistic function
Count	Poisson	identity	Mean count $\lambda_{ki} = \exp(\sum_{j \in C_{ki}} y_{kj} a_{ij} + o_i)$
Count	Poisson	logarithm	Mean count $\lambda_{ki} = (\sum_{j \in C_{ki}} y_{kj} a_{ij} + o_i)$
Count	Negative binomial	identity	Mean count $p_{ki} = \exp(\sum_{j \in C_{ki}} y_{kj} a_{ij} + o_i)$

Table 1: Natural parameter mapping to the expression of the mean for common distributions used for presence/absence or count data, for different choices of the link function.

Association type	Support	Recall (%)	Precision (%)	F1-score (%)
Neutral	[12, 380]	60.75	98.64	74.50
Negative	[2, 20]	72.00	34.02	41.23
Positive	[2, 20]	77.60	17.60	26.72
Averages	-	62.45	94.71	73.09

Table 2: Association classification performances and class supports (number of true associations of each class). The averages are obtained by weighting the score of each association type by its support.