Predicting Links in Plant-Pollinator Interaction Networks Using Latent Factor Models with Implicit Feedback

Eugene Seo,[1] **Rebecca A. Hutchinson**[1,2]

[1]School of Electrical Engineering and Computer Science; [2]Dept. of Fisheries and Wildlife
Oregon State University
Corvallis, OR 97331
{seoe,rah}@oregonstate.edu

Abstract

Plant-pollinator interaction networks are bipartite networks representing the mutualistic interactions between a set of plant species and a set of pollinator species. Data on these networks are collected by field biologists, who count visits from pollinators to flowers. Ecologists study the structure and function of these networks for scientific, conservation, and agricultural purposes. However, little research has been done to understand the underlying mechanisms that determine pairwise interactions or to predict new links from networks describing the species community. This paper explores the use of latent factor models to predict interactions that will occur in new contexts (e.g. a different distribution of the set of plant species) based on an observed network. The analysis draws on algorithms and evaluation strategies developed for recommendation systems and introduces them to this new domain. The matrix factorization methods compare favorably against several baselines on a pollination dataset collected in montane meadows over several years. Incorporating both positive and negative implicit feedback into the matrix factorization methods is particularly promising.

Introduction

Animal pollinators provide essential ecosystem services for agriculture (Bailes et al. 2015). While concerns about honeybee colony collapse receive more media attention, many wild pollinator species also play an important role in this system, and their populations are declining (Tylianakis 2013). Pollination systems can be represented by bipartite networks of interactions between species of two types: plants and pollinators. Nodes in the network correspond to species, and the links between them indicate how frequently the pollinator species visits the plant species; a small example is presented in Figure 1[1]. The scientific questions about pollination systems revolve around the topological structure of these networks and their resilience to perturbations like species extinctions, invasive species, and temporal shifts in flowering and flying periods due to climate change (Pawar 2014; Burkle and Alarcón 2011; Olesen et al. 2007; Memmott, Waser, and Price 2004).

[1]Photos from www.discoverlife.org, Wikipedia, Hartmut Wisch, Greg Lief, www.joycreek.com, science.halleyhosting.com, Robert Pavlis.

Data on pollination networks is collected through visual field observations by recording every visit from a pollinator to a plant as a link between the species. These records are necessarily incomplete since it is impossible to detect every interaction that occurs. Current work studying plant-pollinator systems relies primarily on descriptive statistics of the structure of observed networks (e.g. connectance, nestedness (Pawar 2014; Bascompte et al. 2003)), despite evidence that these statistics are affected by incomplete sampling of the networks (Fründ, Mccann, and Williams 2016; Nielsen and Bascompte 2007). Models that estimate the true network from imperfect observations could allow scientists a more complete view of these systems in studies of the structure and function of pollination networks.

In this paper, we study a set of models for predicting the interactions in pollination networks. We explore network models developed in the context of recommendation systems, which are generally formulated with users and items. One of the most popular examples is the Netflix challenge dataset, which shows users' ratings for movies that they have watched (Bell and Koren 2007). In the analogy to pollination networks, the users are the pollinator species (instead of movie watchers, shoppers, etc.) and the items are the plant species (instead of movies, products, etc.). Since observers are unable to detect every interaction in the field, we seek to 'recommend' which plants may be of interest to which pollinators.

Analyses of movie rating matrices and plant-pollinator interaction networks are similar in that they both aim to capture underlying mechanisms such as preferences. These preferences can determine relationships between two disjoint sets of entities. Both datasets can be represented as bipartite graphs with links between two types of nodes. Based on the links of the graph, both domains try to predict the unobserved hidden links and the most likely links to appear by solving a recommendation problem.

Although movie recommendation and pollination domains share some structure and goals, there are also important differences. First, the movie ratings are provided by active users providing *explicit feedback*, while the pollination data are recorded by observers. As such, the pollination data are more similar to recommendation system domains that use *implicit feedback*, like the amount of time watching a television program or the number of times a crowdsourc-
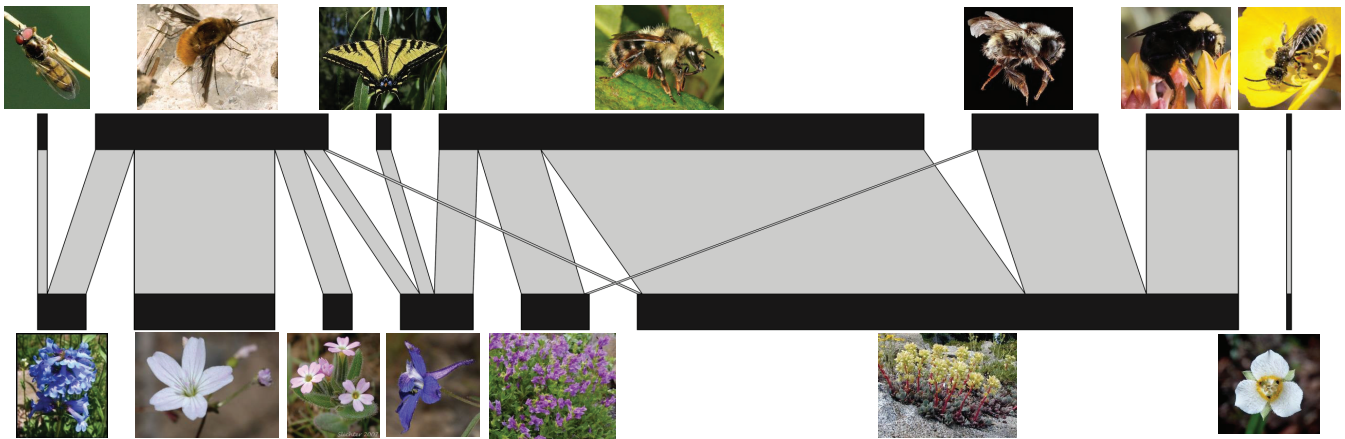
Figure 1: A subset of a pollination network between seven pollinator species and seven plant species. The width of the links between species indicates the frequency of their interactions.

ing task was performed (Hu, Koren, and Volinsky 2008; Lin, Kamar, and Horvitz 2014). Second, movie ratings (and television watching and crowdsourcing tasks) pertain to a single user and single item, while the pollination data accumulates interactions between individual plants and individual pollinators during the survey period. This means that dynamic environmental factors can influence the pollination data, like the abundance of each species during different survey periods. Finally, the pollination data are usually collected for hundreds of species at a time, whereas other recommendation domains often consider data from tens of thousands of users or more.

There are two standard approaches in recommendation problems: content-based approaches and collaborative filtering approaches. Collaborative filtering is better studied than the content-based approach because it can discover hidden relationships between users or items with no additional information beyond interaction values such as ratings. A classic method of collaborative filtering is a nearest-neighbor technique that finds similar users or items based on their interactions and makes a recommendation based on members of the same group. However, matrix factorization techniques have been shown to outperform the classic nearest-neighbor techniques for recommendations (Koren, Bell, and Volinsky 2009; Yang et al. 2014; Hernando, Bobadilla, and Ortega 2016). Matrix factorization decomposes the observed data matrix into latent factors that can characterize both items and users.

This paper contributes the first analysis of a pollination network with matrix factorization techniques, which is conducted on a freely available dataset (Josnes and Pfeiffer 2017) and compares against several alternative approaches for link prediction. Unlike other interaction networks (e.g social networks and protein-protein networks) which have been intensively studied using various network modeling approaches (Rho 2005; Li, Fang, and Sheng 2017), pollination networks have received little attention in computer science. Below, we describe promising results for matrix factorization methods in pollination systems, despite having substan-

tially fewer data than the traditional settings.

Motivating Dataset

Data on pollination networks in montane meadows have been collected at the H.J. Andrews Forest in the Cascade Mountains of Oregon each summer since 2011 (Josnes and Pfeiffer 2017). Observers visit 12 meadows about five times each over the course of the summer. We call each visit to each meadow a *meadow-watch*. On each meadow-watch, the observers record every visit of a pollinator to a flowering plant during each of 15 1-minute periods at each of ten 3 × 3 meter plots in the meadow. Over 112 plant species and 519 pollinator species have been recorded since data collection began. We form the pollination networks by counting the number of meadow-watches in which an interaction occurred. Only 7-10% of the potential interactions are noted in any given year. Table 1 shows the species and interaction counts by year.

Year	Num. Species		Num. Interactions	
	Pollinators	Plants	Common	Unique
2011	137	56	456	273
2012	166	58	446	266
2013	173	63	517	424
2014	113	53	362	146
2015	116	36	242	104

Table 1: Data summaries from the plant-pollinator interactions observed at the H.J. Andrews in each year. Columns show the number of pollinators observed, the number of plants observed, the number of interactions observed that were common to all years, and the number of interactions observed that were unique to this year.

Before conducting this survey of interactions, they count all flowering plant species present in each plot. This information reveals the availability of the plants to be selected by the pollinators. Below, we use these data to inform the implicit feedback models.

Problem Statement

We address the problem setting in which we are given observation frequencies for the links in a bipartite network. In the pollination domain, differences in species abundance and their distribution in time and space affect the observed interaction network substantially, because the pollinators' innate preferences will manifest differently in the data based on the flowers that are available. The species abundances and distribution define a particular context in which the interaction network was observed. Our goal is to predict which links will occur in a new context; that is, which interactions will happen in a new environment with different abundances of flowers from the training network. We are particularly interested in methods that estimate factors driving the pollinators' preferences in order to make these predictions, as the preferences estimates are interesting to ecologists. Note that both sets of entities in the bipartite network are the same in both the training and testing contexts; that is, we do not address the 'cold-start problem' (Koren, Bell, and Volinsky 2009) of generating predictions for brand new plants or pollinators. We study this task by predicting the interactions that occurred in one year of the pollination data from the other years, using only the species that were common across years.

Methods

Classic Matrix Factorization (MF)

Matrix factorization is a latent factor model that discovers hidden factors determining relationships between two entities. Let R be an $m \times n$ data matrix containing the interaction frequencies between m pollinators (users) and n plants (items), and let P be a binary version of this matrix indicating whether the interaction was ever observed or not. This model decomposes P into a product of two matrices, U and V. U is an $m \times k$ matrix describing the pollinators in the latent k-dimensional space, and V is a $k \times n$ matrix describing the plants in the same k-dimensional latent trait space. To learn the decomposed matrices, this model minimizes the error between original data and the approximated data from a product of two matrices and applies regularization to prevent overfitting. The loss function is

$$L(U, V) = \sum_{ij} (p_{ij} - u_i^T v_j)^2 + \lambda(\|U\|_F^2 + \|V\|_F^2),$$

where p_{ij} indicates whether pollinator i interacted with plant j, u_i is pollinator i's trait vector, v_j is plant j's trait vector, and λ is a regularization parameter. We optimized this loss function using gradient descent with alternating updates to U and V for 20 iterations.

Matrix Factorization with Implicit Feedback (IFMF)

The implicit feedback matrix factorization model (IFMF) addresses the case in which explicit feedback such as a rating is not available (Hu, Koren, and Volinsky 2008). This method uses other information about users, such as the frequency of watching TV programs or the time spent on web sites, as evidence of their preferences. IFMF utilizes this implicit feedback as a confidence for the binary preference about whether a user likes an item or not rather than considering it as explicit feedback. It introduces weights, w_{ij}, representing the confidence, defined as $w_{ij} = 1 + \alpha c_{ij}$, where c_{ij} is the implicit feedback and α is a tuning parameter. IFMF minimizes the loss function

$$L(U, V) = \sum_{ij} w_{ij}(p_{ij} - u_i^T v_j)^2 + \lambda(\|U\|_F^2 + \|V\|_F^2),$$

where p_{ij} indicates whether a link was observed. Again, we used 20 iterations of alternation gradient descent updates.

Matrix Factorization with Extended Implicit Feedback (IFMF2)

IFMF only incorporates positive implicit feedback. More recent work inspired by the crowdsourcing domain extends this approach to IFMF2 by incorporating negative implicit feedback as well (Lin, Kamar, and Horvitz 2014). Just as the positive implicit feedback measures frequent interactions as indications of preference, IFMF2 infers negative feedback when an item was plentiful but unused. In the domain of task recommendation for crowdsourcing, the number of completed tasks served as positive implicit feedback (c_{ij}) and the availability of tasks that a user has *not* worked served as negative implicit feedback, termed d_{ij}. Logistic functions,

$$f(x) = \frac{1}{1 + \exp(\alpha_1(\log(x) + \beta_1))}$$

and

$$g(x) = \frac{1}{1 + \exp(\alpha_2(\log(x) + \beta_2))}$$

normalize the ranges of the positive and negative implicit feedback to similar scales, and the weights in IFMF2 are defined as $w_{ij} = f(c_{ij})$ for observed links ($p_{ij} = 1$) and $w_{ij} = g(d_{ij})$ for unobserved links ($p_{ij} = 0$). The loss function now uses this new definition for the weights:

$$L(U, V) = \sum_{ij} w_{ij}(p_{ij} - u_i^T v_j)^2 + \lambda(\|U\|_F^2 + \|V\|_F^2).$$

We optimized this loss function with the same gradient descent procedure as in MF and IFMF.

Baselines

We compare the matrix factorization approaches to four baselines: Popularity, Availability, User-Nearest-Neighbors, and Item-Nearest-Neighbors.

The Popularity baseline determines all pollinators' preferences for the plants based on the popularity of plants in the training set. The popularity is measured as the total number of interactions the plant has from all pollinator species. This baseline assumes that all pollinators have the same preferences, which are reflected in the plants' visit frequencies. If the popular plants in the training set are also available and popular in the test set, this baseline can perform well. The Availability baseline is based on the abundance of the plants in the test set. This baseline assumes that pollinators visit plants in proportion to their availability.

Nearest-neighbor algorithms are a basic family of approaches for collaborative filtering. This technique aims to find similar groups of either users or items and make predictions for unrated items based on the groups. The User-Nearest-Neighbors (UserNN) baseline is also based on the idea that pollinators' preferences will be reflected in the behavior of its nearest neighbors. The similarity of two pollinators is calculated by measuring the cosine similarity between their interaction frequencies as

$$s_{i_1 i_2} = \frac{r_{i_1}^T r_{i_2}}{\|r_{i_1}\|\|r_{i_2}\|},$$

where r_i is ith row of R. We predict the pollinator i's preference for the flower j as

$$\hat{r}_{ij} = \frac{\sum_{u \in S(i;j)} s_{iu} r_{uj}}{\sum_{u \in S(i;j)} s_{iu}},$$

where $S(i;j)$ is the set of neighbors of pollinator i that have interacted with plant j. The Item-Nearest-Neighbors (ItemNN) baseline is analogous for plant-based similarity.

Experiments
Defining implicit feedback for pollination networks

In recommendation systems for TV watching, the amount of time spent watching a program served as positive implicit feedback (Hu, Koren, and Volinsky 2008). The number of tasks completed was treated as positive implicit feedback for a crowdsourcing recommendation system (Lin, Kamar, and Horvitz 2014). In the pollination domain, the positive implicit feedback is based on the frequency with which each pollinator visited each plant. If a pollinator has interacted with a plant frequently, we are more confident that the pollinator has a strong preference for this plant and that the pairwise interaction will occur in a new context.

The creators of IFMF2 used the availability of tasks that users have not completed as negative implicit feedback, measured by the total number of times each task was completed (summed over all users) (Lin, Kamar, and Horvitz 2014). In this domain, the negative implicit feedback is based on the availability of the plants, as determined by the flower surveys that preceded the interaction surveys. Confidence for a negative preference increases if a plant that a pollinator never interacted with was highly available. Here and in the prior work on IFMF2, all plants (tasks) are equally available to all pollinators (users). The availability takes the form of a vector with an entry for each plant, so we refer to this as the *vector-based availability* case, or *A.vec* for short.

The vector-based availability assumes that all pollinators have equal access to the plants, but this assumption is faulty. A plant that flowers abundantly in early summer may appear as highly available in the vector, but for pollinators that do not emerge until late summer, that plant may be completely unavailable. To account for mismatches between pollinators and plants in space and time, we consider a second approach, which we call *matrix-based availability*, or *A.mat* for short. In this case, the presence of the plant in a meadow-watch is only counted if the pollinator was also present in the same meadow-watch. Since the pollinators are not surveyed separately, we get this information from the interaction survey and determine that a pollinator is present if it interacted with any plant species on that meadow-watch.

Performance metrics

The network models make predictions about the links in the network. Two possible ways to evaluate the predictive performance of these models are accuracy in distinguishing present versus absent links or root-mean-squared error (RMSE) between predicted and actual link weights. However, both of these evaluation metrics are problematic in this domain. Accuracy obscures the differences between links that are frequent and rare. RMSE overemphasizes the count values on the links, which are observed imperfectly and have widely ranging values (e.g. 0-60). Instead of link occurrence or weights, we focus on the ranking of the interaction partners for each pollinator to evaluate the methods. This ranking evaluation measures the correctness of predictions in terms of relative likelihood of interactions between plant species and pollinator species. We use two metrics to evaluate these predicted ranks: mean percentile rank (MPR), and precision-recall (PR) curves. MPR and PR curves have been used to evaluate these methods in other domains (Lin, Kamar, and Horvitz 2014).

Both of these methods require the interactions to be ranked in order to compute the metrics. Before ranking the interactions, we calculate the weights of the links by multiplying the predicted preferences by availability to get the total number of interactions expected. We used both versions of the availability measurements described above (*A.vec* and *A.mat*). Then, the expected totals produce the final ranking of plants for each pollinator.

Mean percentile rank is calculated as

$$\frac{\sum_{ij} c_{ij} \rho_{ij}}{\sum_{ij} c_{ij}},$$

where ρ is the percentile rank of plant j by pollinator i. If ties occur, we average their ranks to penalize the repetition. The percentile rank of the most favored plant is 0 and the least favored is 100, so a lower MPR is better. We compute lower bounds on the MPR using the percentile rankings from the test set. Since MPR weights the percentile rankings by the link frequencies in the test set, it is sensitive to the scale of those frequencies. Small errors in the ranks of plants that are interacted with much more than their counterparts are penalized more heavily than similar errors on plants that are interacted with infrequently.

Precision-recall curves provide a complementary metric that focuses less on getting the ranking exactly correct and more on the approximate ordering. To construct them, we retrieve the top N pairwise interactions based on the predicted rankings, varying N from 1 to the total number of plants to see the balance of precision and recall in a relative curve. In the case that the rankings have duplicated values, we pick one of them randomly to maintain same number of top N recommendation for all methods.

All Interactions	2011 (LB: 3.65)		2012 (LB: 3.96)		2013 (LB: 4.44)		2014 (LB: 5.12)		2015 (LB: 5.25)	
	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat
Popularity	**10.29**	**9.18**	16.32	**13.86**	**17.12**	**15.90**	**11.63**	**11.17**	15.73	**14.97**
Availability	27.54	21.45	29.98	22.73	29.34	25.43	23.32	22.07	26.42	23.74
ItemNN	30.70	25.38	33.05	26.46	32.24	27.12	23.54	22.48	24.11	21.96
UserNN	14.25	12.10	18.73	15.25	19.45	17.06	15.21	14.36	21.87	19.89
MF	17.39	13.77	21.11	17.18	23.05	21.03	17.58	15.83	23.93	21.76
IFMF	11.45	11.87	**15.44**	17.13	17.48	19.91	12.24	15.76	**13.38**	19.68
IFMF2	15.36	9.33	20.20	15.29	20.12	19.50	16.47	13.46	19.68	16.37

New Interactions	2011 (LB: 2.47)		2012 (LB: 1.46)		2013 (LB: 1.99)		2014 (LB: 3.05)		2015 (LB: 4.43)	
	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat	A.vec	A.mat
Popularity	23.97	21.02	**22.28**	16.52	29.79	23.16	21.74	19.04	29.49	21.81
Availability	26.74	21.21	34.12	20.11	35.88	25.65	22.61	23.92	40.11	23.21
ItemNN	27.24	21.60	33.23	19.79	33.61	23.65	22.84	20.57	40.93	26.27
UserNN	28.12	22.53	24.25	17.61	31.65	22.55	23.10	20.22	**27.33**	**18.86**
MF	23.55	19.17	29.03	18.76	**29.68**	22.66	22.20	20.18	42.11	28.27
IFMF	23.48	17.44	24.38	16.67	30.58	22.18	**20.84**	22.85	29.24	22.77
IFMF2	**21.76**	**17.26**	23.47	**13.10**	30.23	**19.58**	22.35	**18.40**	37.06	21.53

Table 2: MPR of five testing years (2011-2015) for each model. The top half shows results for predicting all of the interactions in the test set, and the bottom half shows results for predicting only the new interactions in the test set. In each set, the left column considers the availability of interactions based on a vector of plant availabilities (A.vec), whereas the right column considers the availability of each interaction based on the availability of both the plant and the pollinator (A.mat).

Parameter tuning

We formed five test sets, each comprising the data from one year (2011-2015). We performed cross-validation within the training set to tune parameters, again partitioning the data into folds based on year (i.e. four-fold cross-validation). We tuned k among the values $\{2, 5, 10, 15\}$ for all matrix factorization approaches, the IFMF parameter α among the values $\{1, 10, 20\}$, and the logistic functions parameters $\alpha_1, \alpha_2, \beta_1, \beta_2$ among the values from -10 to -1 in increments of 0.5. We chose the optimal parameters based on MPR values.

We evaluated performance on the test set based on two different sets of interactions. In the first case, we used *all* interactions to compute the performance metrics, even if they had already appeared in the training set. While this case gives credit to the algorithms for predicting interactions that have already been seen, the ability to predict these interactions in the new context holds interest in the ecological setting. In the second case, we used only interactions that were *new* to the test set to compute the metrics. This case focuses on the ability to predict novel interactions between species that did not interact in the training set.

Results

When predicting all interactions in the test sets, the Popularity baseline performed well (Table 2). Only in two cases did another method (IFMF) beat this baseline; in at least one of these cases, a plant that was unpopular in the training set was extraordinarily popular in the test year. Recall that MPR is highly influenced by the most frequent interactions.

These are the same interactions that are well-represented by the Popularity baseline. When predicting only the interactions that are new to the test set, the situation is reversed. Only in one case does the Popularity baseline perform best; in all other cases, a matrix factorization method with implicit feedback (usually IFMF2) is better. The new interactions in the test sets are rarer; on average, an interaction common to all years appeared in 7.1 meadow-watches, but the interactions that only occurred in one year appeared in 3.8 meadow-watches. When focusing on the novel interactions, the Popularity baseline is too general, but the latent traits learned by the matrix factorization methods are informative. In most cases, the matrix factorization methods outperformed the Availability baseline and the neighborhood-based methods. The Availability baseline fails on specialist species, which do not have many interactions even when they are highly abundant. The neighborhood-based methods suffered from the sparsity of the interactions. On average, pollinators do not interact with many plants, so it is difficult to find good neighbors and the predictions are very sparse. The test year 2015 was the exception to some of these trends. 2015 was an unusual year due to drought, and the 2015 network is the smallest and most densely connected of the years.

With a few exceptions, encoding the availability information as a matrix (*A.mat*) is better than as a vector (*A.vec*). This is intuitive, since the matrix form takes into account the presence of both the pollinator and the plant to assess availability of the interaction to occur, whereas the vector version only considers the plants. Especially when predicting
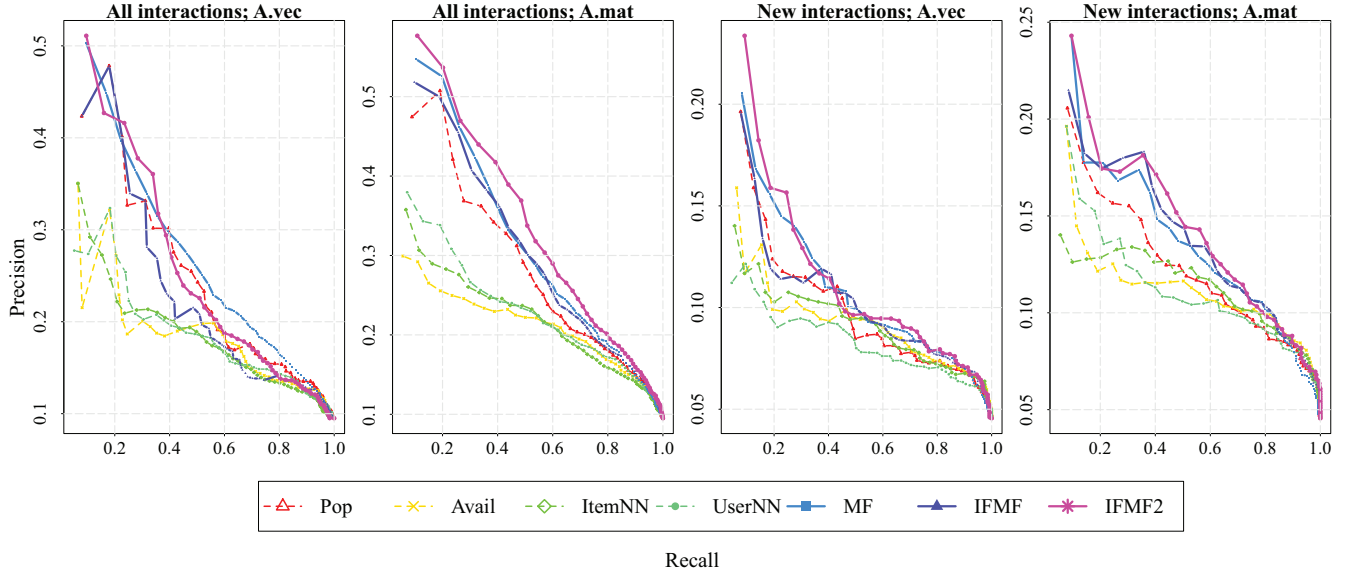
Figure 2: Precision-Recall curves for all methods in test year 2011. Plots show results for predicting all interactions vs. only new interactions and considering availability based on plants only (A.vec) vs. plants and pollinators (A.mat).

new interactions, the IFMF2 method performs the best with *A.mat* among all methods except in test year 2015. Thus, it appears that negative implicit feedback and the personalized availability information improve performance.

Whereas MPR focuses heavily on ranking the tasks perfectly, the PR curves inform the models' ability to predict more frequent interactions above less frequent interactions more generally. Figure 2 shows a representative example of the PR curves for the 2011 test year; plots for the other years are in the supplemental material. While the Popularity baseline remains strong, the PR curves show more differences between the matrix factorization methods and Popularity, even when predicting all interactions in the test set instead of just the novel ones. This evidence further supports the explanation that the MPR is strongly influenced by the most frequent interactions. The gap between Popularity baseline and the matrix factorization methods is getting also more distinctive when predicting with *A.mat* over *A.vec*, as with MPR. This implies that the traits learned from the matrix factorization methods are close to the hidden factors that determine the plant-pollinator interactions regardless of dynamic changes in flower abundance. Again, the neighborhood-based methods and the Availability baseline are not competitive overall.

In addition to evaluating predictive performance, we visualized the latent factors learned by an IFMF2 analysis trained on all five years of data. The tuning parameters were chosen using leave-one-year-out cross-validation on the entire dataset. Since the number of latent factors was chosen to be $k = 2$, we can plot the species in the latent feature space directly (i.e. without an additional dimensionality reduction step). The results are plotted in Figure 3. We consulted with a pollination expert to interpret the plots. For the pollinators

(U; top plot), the first latent factor (horizontal axis) places rare species on the left and more common species on the right. The second latent factor (vertical axis) shows a subtle trend from generalist species at the top to specialist species at the bottom. For the plants (V; bottom plot), the first latent factor (horizontal axis) sorts species by the degree of access to their flowers, with flowers that exclude some pollinators based on their shape or size on the left and flowers that almost any pollinator can access on the right. Among the flowers with exclusionary shapes, the second factor (vertical axis) was correlated with the degree of the exclusion. The flowers lower on the axis excluded more pollinator species than those higher on the axis. Further analysis of the species in the latent space is available in the supplemental material, including automated clustering of the species, expert descriptions of the clusters, and correlations of the latent factors with species characteristics to aid interpretation.

Discussion

The experiments presented above show promising results for matrix factorization techniques in a new domain. The pollination domain is characterized by smaller, more densely-connected networks than applications in movie or television recommendation and crowdsourcing. The pollination networks also show a strong impact of species availability on the interactions that are observed. In particular, not all items in this domain are equally available to all users, due to variation in phenology, morphology, and abundance across species, time, and space. This leads to greater interest in predicting all interactions rather than just new ones, since the test interactions may occur in a substantially different context than the training interactions. We found that matrix fac-

Pollinator species in the latent factor space

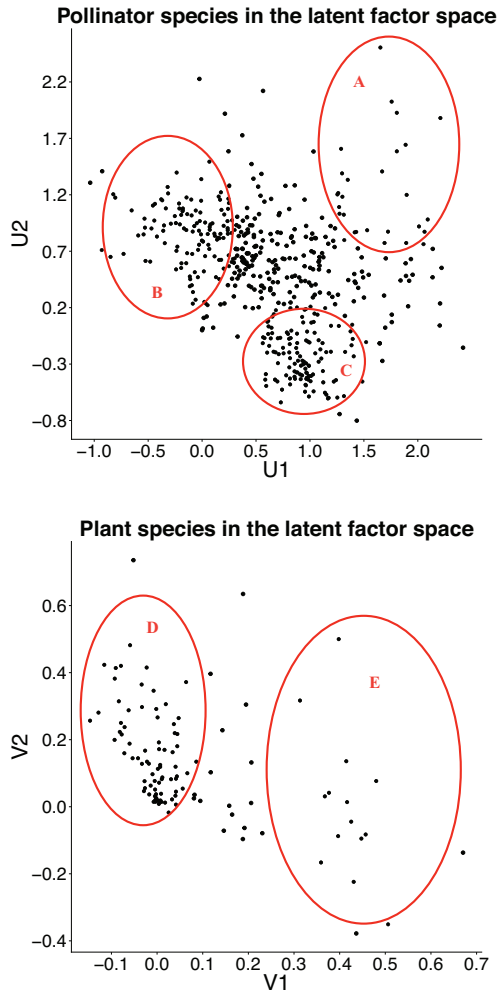**Plant species in the latent factor space**

Figure 3: Species plotted in the latent factor space learned by IFMF2 with $k = 2$ factors. Circles indicate trends noted by a pollination expert: A) common pollinators that visit many plant species; B) rare pollinators; C) common pollinators that visit a small subset of species; D) plants that exclude some pollinators through size or shape mismatches; E) plants whose flowers are open to almost any pollinator.

torization methods help to predict the interactions that will occur in these new contexts, especially when focusing on the novel interactions. The implicit feedback models that make use of availability information perform better than the alternatives in terms of predicting the interactions that will be observed in a new context with different species abundances.

Our results showed that it is helpful to incorporate implicit feedback as a matrix that contains availability information specific to each user-item pair rather than a vector that reflects the items more generally. However, an availability matrix must be constructed for the new context into which we wish to predict. This presumes that this information can be gathered for the test set. In this case, we consider a pollinator to be available if it interacted with any plant in the test year. In contrast, the vector-based availability only re-

quires information from the flower survey for the test year. The relevance of these results will vary based on the kind of information that can be gathered about availability in the new context. If a natural resource manager seeks advice on what to plant in a meadow to bolster certain pollinator populations, they may only have availability information as proposed flower abundances. On the other hand, if they are introducing hives or nests for some pollinator species, they may have more information about pollinator availability for creating a matrix. Similarly, their interest in predicting all interactions versus only novel interactions may vary based on the particular application scenario.

While our results are promising, the performance gains are modest in some cases, and the lower bounds on MPR indicate that there is room to improve the predictions in future work. One area for improvement may be parameter tuning, as the tuning parameters were chosen based on MPR, and we observed some sensitivity to these settings; tuning for precision and recall may be better. We plan to investigate sensitivity to the logistic function parameters in more detail and to explore alternate strategies for weighting the interactions in the loss functions.

We are also interested in pursuing alternatives to matrix factorization for this problem. Prior work in social network analysis has incorporated known traits as well as latent factors (Hoff 2009), which could leverage information about the species in the network. Work in nonparametric Bayesian models of bipartite networks has incorporated phylogenetic priors (Miller, Griffiths, and Jordan 2008), and work is underway to construct a phylogeny of the species in this dataset. Finally, since these networks are observed imperfectly in the field, we are interested in models of missingness for these data, which has also been studied in the collaborative filtering context (Marlin et al. 2007).

Conclusion

This paper introduced pollination networks as a new domain for models and algorithms developed for recommendation systems. This domain differs from traditional applications in that the networks are smaller, denser, and sensitive to the variable availability of the plants and pollinators to interact. We conclude that matrix factorization methods are a promising family of methods for predicting interactions in pollination networks. In particular, methods that incorporate both positive and negative implicit feedback about the availability of the plant and pollinator species to each other are helpful for predicting novel interactions in a new context. These methods may be applicable in other domains in which interactions are recorded at the level of groups (e.g. species) rather than individuals and the sizes of the groups are variable, such as host-parasite networks and food webs.

References

Bailes, E. J.; Ollerton, J.; Pattrick, J. G.; and Glover, B. J. 2015. How can an understanding of plant-pollinator interactions contribute to global food security? *Current Opinion in Plant Biology* 26:72–79.

Bascompte, J.; Jordano, P.; Melián, C. J.; Olesen, J. M.; Melian, C. J.; and Olesen, J. M. 2003. The nested assembly of plant-animal mutualistic networks. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 100, 9383–9387.

Bell, R. M., and Koren, Y. 2007. Lessons from the netflix prize challenge. *ACM SIGKDD Explorations Newsletter* 9(2):75–79.

Burkle, L. A., and Alarcón, R. 2011. The future of plant-pollinator diversity: Understanding interaction networks acrosss time, space, and global change. *American Journal of Botany* 98(3):528–538.

Fründ, J.; Mccann, K. S.; and Williams, N. M. 2016. Sampling bias is a challenge for quantifying specialization and network structure: Lessons from a quantitative niche model. *Oikos* 125(4):502–513.

Hernando, A.; Bobadilla, J.; and Ortega, F. 2016. A non negative matrix factorization for collaborative filtering recommender systems based on a bayesian probabilistic model. *Knowledge-Based Systems* 97:188–202.

Hoff, P. 2009. Multiplicative latent factor models for description and prediction of social networks. *Computational and Mathematical Organization Theory* 15(4):261–272.

Hu, Y.; Koren, Y.; and Volinsky, C. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the Eighth IEEE International Conference on Data Mining*, 263–272.

Josnes, J. A., and Pfeiffer, V. W. 2017. Plant Pollinator data at HJ Andrews Experimental Forest, 2011 to present. Long-Term Ecological Research. Forest Science Data Bank, Corvallis, OR. [Database]. Available at http://andlter.forestry.oregonstate.edu/data/abstract.aspx?dbcode=SA026.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37.

Li, Z. L.; Fang, X.; and Sheng, O. R. L. 2017. A Survey of Link Recommendation for Social Networks: Methods, Theoretical Foundations, and Future Research Directions. *ACM Transactions on Management Information Systems* 9(1):1–26.

Lin, C. H.; Kamar, E.; and Horvitz, E. 2014. Signals in the Silence: Models of Implicit Feedback in a Recommendation System for Crowdsourcing. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 908–914.

Marlin, B. M.; Zemel, R. S.; Roweis, S.; and Slaney, M. 2007. Collaborative Filtering and the Missing at Random Assumption. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, 9.

Memmott, J.; Waser, N. M.; and Price, M. V. 2004. Tolerance of pollination networks to species extinctions. In *Proceedings of the Royal Society of London B: Biological Sciences*, volume 271, 2605–2611.

Miller, K. T.; Griffiths, T. L.; and Jordan, M. I. 2008. The phylogenetic indian buffet process: A non-exchangeable nonparametric prior for latent features. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, 403–410.

Nielsen, A., and Bascompte, J. 2007. Ecological networks, nestedness and sampling effort. *Journal of Ecology* 95(5):1134–1141.

Olesen, J. M.; Bascompte, J.; Dupont, Y. L.; and Jordano, P. 2007. The modularity of pollination networks. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 104, 19891–19896.

Pawar, S. 2014. Why are plant-pollinator networks nested? *Science* 345(6195):383–383.

2005. Probabilistic model of the human protein-protein interaction network. *Nature Biotechnology* 23(8):951–959.

Tylianakis, J. M. 2013. The Global Plight of Pollinators. *Science* 339(6127):1532–1533.

Yang, X.; Guo, Y.; Liu, Y.; and Steck, H. 2014. A survey of collaborative filtering based social recommender systems. *Computer Communications* 41:1–10.