

Identifying a Ranking of Plant Preferences for a Pollinator

by
Sneha Krishna Kumaran

A PROJECT

submitted to

Oregon State University
University Honors College

in partial fulfillment of
the requirements for the
degree of

Honors Baccalaureate of Science in Computer Science
(Honors Associate)

Presented May 11, 2015
Commencement June 2015

AN ABSTRACT OF THE THESIS OF

Sneha Krishna for the degree of Honors Baccalaureate of Science in Computer Science presented on May 11, 2015. Title: Identifying a Ranking of Plant Preferences for a Pollinator.

Abstract approved: _____
Rebecca Hutchinson

Pollinators are an integral part of agriculture and the ecosystem. However, due to changing land use, populations of wild pollinators are decreasing and plant distributions are changing all around the world. To understand how plant-pollinator networks will adapt over time, we would like to understand how pollinators choose flowers to visit. We will model a pollinator's interaction with plant species in two ways: first using a probabilistic multinomial approach to fit a preference score to each plant and second to explain our findings from the multinomial model using the traits of the flowers themselves. Our findings show that a model with preferences performs better than a model which does not have preferences. While this model shows potential in finding plant preferences, it does not fully explain the distribution of plant-pollinator interactions. To try to explain the interactions more fully, we incorporated the traits of the plants into the score of the plant. We found that the traits do have some effect on the score of the plant, but again do not fully explain the interactions in this particular model.

Key Words: Plant-pollinator networks, machine learning, computational ecology, pollinator behavior

Corresponding e-mail address: sneha.krishna@gmail.com

©Copyright by Sneha Krishna
May 11, 2015
All Rights Reserved

Identifying a Ranking of Plant Preferences for a Pollinator

by
Sneha Krishna Kumaran

A PROJECT

submitted to

Oregon State University
University Honors College

in partial fulfillment of
the requirements for the
degree of

Honors Baccalaureate of Science in Computer Science
(Honors Associate)

Presented May 11, 2015
Commencement June 2015

Honors Baccalaureate of Science in Computer Science project of Sneha Krishna Kumaran presented on May 11, 2015

APPROVED:

Rebecca Hutchinson, Mentor, representing Computer Science

Thomas Dietterich, Mentor, representing Computer Science

Julia Jones, Committee Member, representing Earth, Ocean, and Atmospheric Sciences

Toni Doolen, Dean, University Honors College

I understand that my project will become part of the permanent collection of Oregon State University, University Honors College. My signature below authorizes release of my project to any reader upon request.

Sneha Krishna Kumaran, Author

Acknowledgement:

I would like to thank my committee Dr. Rebecca Hutchinson, Prof. Tom Dietterich and, Prof. Julia Jones for taking the time to evaluate my thesis and. I would also like to thank the Eco-Informatics Summer Institute for allowing me to use the data collected over the summers and Dr. Andrew Moldenke for his insight into the behavior of pollinators.

Table of Contents

| | |
|--|----|
| Introduction..... | 1 |
| Motivation..... | 1 |
| Objective..... | 1 |
| Background Information..... | 2 |
| Materials and Methods..... | 6 |
| Field Study Data | 6 |
| Models | 7 |
| Multinomial Model | 7 |
| Traits Model..... | 9 |
| Simulated Data Generation..... | 11 |
| Bootstrapping for Confidence Intervals..... | 12 |
| Evaluation | 12 |
| Goodness of Fit—Chi-Squared Test..... | 12 |
| Correlations..... | 13 |
| Likelihood Ratio Test | 14 |
| Results..... | 15 |
| Multinomial Model | 15 |
| Using Simulated Data | 15 |
| Experiments on the HJ Andrews Data: | 18 |
| Bootstrapping from the HJ Andrews Field Data..... | 22 |
| Traits Model..... | 28 |
| Preliminary testing | 28 |
| Simulated Data Experiments..... | 31 |
| Experiments on the HJ Andrews Data | 33 |
| Bootstrapping from the HJ Andrews Dataset | 34 |
| Comparing the Multinomial Model with the Traits Model..... | 35 |

| | |
|---|----|
| Discussion | 36 |
| Conclusions | 36 |
| Next Steps | 37 |
| Bibliography | 39 |
| Appendix A: The Binary Preference Model | 41 |
| A Multinomial Preference Model | 41 |
| Multinomial Model | 41 |
| Binary Model | 42 |
| Synthetic Data Analysis | 43 |
| The Algorithm | 43 |
| The Experiments | 44 |
| Conclusions | 45 |
| Appendix B: The Burstiness Model | 46 |
| The Model | 46 |
| A Proposed Algorithm | 47 |

Introduction

Motivation

Environmentalists and policy makers are worried about the decline of insect pollinators (such as bumblebees, moths, and butterflies) (Tylianakis, 2013), which are important to both the agricultural industry and many ecosystems (Tylianakis, 2013; Aizen, et al., 2009). These pollinators are disappearing, and it is neither a local phenomenon nor a single species. Populations of wild pollinators are decreasing around the world due to changing land use (Fitzpatrick, et al., 2007; Biesmeijer, et al., 2006). Ecologists predict that the loss of pollinators will affect plant biodiversity (Tylianakis, 2013). Pollinators are also important in agriculture because around 30% of the plants we eat require pollination in order to produce fruits and vegetables (Tylianakis, 2013). As wild pollinator populations decrease, we must supplement them with managed honeybee hives. However, this comes with a price as studies have shown that the yield for these plants increases when they are pollinated by wild pollinators (Aizen, et al., 2009).

Experts are also concerned about how plant extinctions will affect insect pollinator populations, and there is evidence that pollinator extinctions are followed by plant extinctions (Biesmeijer, et al., 2006; Aizen, et al., 2012). However, there is also evidence suggesting that pollinators, even ones that have been observed exclusively interacting with only a few plants, can adapt to changing plant species distributions (Tylianakis, 2013). Understanding more about how pollinators choose the flowers they visit will help us predict how plant-pollinator networks will respond to changes in the distribution and abundance of the species within the network.

Objective

The objective of this study is to create a method to identify a ranking of plant preferences for a pollinator given field observations of plant-pollinator interactions and to determine how floral traits affect the preference for each plant. To do this, we developed a probabilistic model to identify the preferences in interaction networks under the assumption that the plant-pollinator interactions can be modelled using only the frequency of the interactions. We used the data collected at the H.J. Andrews Experimental Forest by students from the Eco-Informatics Summer Institute (Pfeiffer, 2011-2014) and explored an extension of this model to determine if the traits of the flowers affect the preferences of a pollinator. In this thesis, we will focus on identifying the preferences of pollinators that interact with a large variety of plant species.

We applied these models to analyze plant-pollinator observations collected in the field. These models take steps towards creating models that will be able to predict plant-pollinator interactions in the case of the extinction of existing plants in a network, the

introduction of invasive species of plants to the area, or the restoration of an extirpated plant to an area.

Background Information

To understand the behavior of pollinators (e.g., to determine their preferences), we must first understand the background of plant-pollinator interactions. The relationship between insect pollinators and flowers is mutualistic, as flowers need to be pollinated and insects use flowers as a source of food, so plants and pollinators co-evolved. To ensure that they can gather enough food, pollinators do not want to compete with other species, so there is a push to focus only one plant species. However, if the pollinator becomes dependent on this one plant species, and this species becomes extinct, the pollinator may also be in danger of going extinct. The same is true for plants and pollination. A *generalist* plant or pollinator is one that interacts with many species of the other group. A *specialist* plant or pollinator only interacts with one or few species of the other group.

There are a few species that are well known to be truly specialist pollinators (such as species in the genus *Dufourea*—each species specializes on a different flower) (Moldenke, 2015). However, much specialization occurs due to the physical characteristics of flowers that limits some pollinators from visiting. For example, some flowers hang suspended which does not allow the pollinator to land on the flower and the pollinator must fly into the flower (Moldenke, 2015). Collecting nectar and pollen in this manner takes energy and certain skills, which not all pollinators have. As another example, some plants have petals closed around the rewards (pollen and nectar), so the pollinators for those plants need to have the skill to open the flower to access the rewards. In cases where unrelated plants share similar pollinators and physical attributes, the shared physical attributes are called *pollination syndromes*. We have some evidence supporting pollination syndromes, but we need more critical examination of this theory (Johnson & Steiner, 2000).

Many plants, such as plants that are highly dispersed and perennials, benefit greatly by being pollinated by specialist pollinators. Plants that are highly dispersed would probably specialize on one pollinator that is loyal to that plant species so that the stigmas are not clogged by pollen of more abundant plants (Johnson & Steiner, 2000). The perennials in this study are highly dispersed, but perennials also live over multiple years, so they have multiple chances for pollination and the opportunity to specialize on just a few loyal insects.

Other plants may benefit from being pollinated by generalist pollinators (Johnson & Steiner, 2000). Annual plants encourage generalist pollinators (Johnson & Steiner, 2000)

because they only have one chance at pollination and attracting as many pollinators as possible means more chances of pollination in that short period of time. Annuals also vary greatly in abundance year to year, so pollinators are unlikely to specialize on these plants (Moldenke, 2015). Abundant plant species and plant species with separate sexes are also predicted to be pollinated by many different pollinators, but especially generalist pollinators (Johnson & Steiner, 2000).

While we have many studies on pollination syndromes and the specialization of plants and pollinators, we must remember that some pollinators could be labeled incorrectly as specialists due to sampling error. If a pollinator is very rare, researchers will observe only a few of the interactions it makes. If these few interactions are with the same plant species, the pollinator might be labelled a specialist without enough evidence (Bluthgen, 2010). There are many methods to determine if a species was present in the field but was not detected, from statistical models (MacKenzie, et al., 2002) to machine learning models (Hutchinson, et al., 2011). With the data collection protocol followed for the data used in this thesis, the previous methods are not needed. However, this shows that the abundance of plants is also critical when trying to model plant-pollinator interactions. If a plant is very abundant, a generalist pollinator might visit this plant even if it prefers a rarer plant much more.

Competition also plays a huge role in field studies like the one used for this study. If a flower is open both at night and in the day, many nocturnal pollinators (such as moths) may visit the flower at night leaving less nectar for the pollinators who visit during the day. Researchers during this field study also noticed that bumblebees and honeybees tend to start their foraging earlier in the morning than solitary bees. This could mean that solitary bees visit flowers with less reward because the nectar from other flowers was already consumed by the social bees (Pfeiffer, 2011-2014). This would cause the solitary bees to behave as if they are specialized in flowers that bumblebees and honeybees do not like.

Other pollinator preferences have also been studied such as flower symmetry, flower color, and *corolla* (petal) shape (Gomez, et al., 2008; Moller & Sorci, 1998; Rodriguez, et al., 2004). These studies show that the corolla shape is a very important factor in the preferences of pollinators. However, corolla shape could be important only due to other related factors. For example, flying and hovering over a flower consumes much of a pollinator's energy. Therefore, corolla shape may only be important because the petals form a platform for the pollinator to land on or allow the pollinator to walk to the next flower. Researchers have also hypothesized that corolla size and shape could give the pollinator a clue as to the reward from visiting the flower. Flower color has been a popular target for pollinator preference studies (especially bees). The theory is that

brightly colored flowers (whites, yellows and oranges), reflect ultraviolet rays making them more visible to insects. Reds and dark colors are theoretically not as visible to pollinators (Moldenke, 2015), but studies of German wildflowers and lowland tropical plants show that flower color by itself is not significantly correlated with pollinator preferences (Johnson & Steiner, 2000).

Evolutionary theory states that if pollinators are specialists, and the plant they depend on becomes extinct, the specialist pollinator will also go extinct (Biesmeijer, et al., 2006; Memmott, et al., 2007; Aizen, et al., 2012). However, some studies suggest that declines of specialist populations are not caused by simple food-plant specialization (Fitzpatrick, et al., 2007) and others find that the specialist pollinators are the ones acquiring new plant interactions (Tylianakis, 2013). Field studies today contain interactions that are not recorded in historical studies and tend to involve species with previously narrow diets (Tylianakis, 2013). This could mean that even specialist pollinators can adapt to human alterations of the landscape, such as introduction of foreign or invasive species. Ultimately, pollinator preference modeling could aid conservation biologists by finding which plant, if any, could be introduced to an area to best help a struggling pollinator species.

Because the quantity and quality of pollination experienced by plants is also declining (Tylianakis, 2013), a few species of generalist pollinators is not enough to keep plant biodiversity intact. For example, the managed honeybee (*Apis mellifera*) populations are now commonly being raised on agricultural fields to supplement pollination normally performed by wild pollinator populations (Tylianakis, 2013; Aizen, et al., 2009). However, agricultural crops that require pollination yield more produce when pollinated by a variety of wild pollinators regardless of the presence of *Apis mellifera* (Aizen, et al., 2009; Tylianakis, 2013). Experts found that honeybees are not equally efficient at pollination of different plant species even though they frequently transport a lot of pollen (Tylianakis, 2013).

Relative to this thesis, there are two relevant approaches to the study of insect pollination: network analysis and modelling and preference modelling. In network analysis, plant-pollinator interactions are modeled as bipartite graphs (Bascompte, et al., 2003). Bipartite graphs have two disjoint sets of nodes and connections only exist between nodes in different sets. No connections occur between two nodes in the same set. In this case, plant species and pollinator species make up the two sets of nodes and the interactions between plants and pollinators are translated to connections. Bipartite networks are used to study plant and pollinator network structure. Research on these networks has shown that in large field studies, generalist pollinators visit nearly all plant species and most specialist pollinators visit plants species that are already visited by many generalist pollinators

(Bascompte, et al., 2003; Moldenke, 1979). Network analysis, while not used in this thesis, could be helpful as we grow our models from looking at a single pollinator species to multiple pollinator species.

Another modelling approach is to model the behavior of pollinators due to apparent preferences. One study examined how pollinator preferences for wild flowers in field margins were affected by agricultural monocultures (Rands & Whitney, 2010). The model created in this study found that pollinators located within agricultural fields preferred the wild flowers in the field margins as the density of the wild flowers increased. This study was based on the idea that pollinators show density-dependent preferences, sometimes even avoiding plants that are unfamiliar (Rands & Whitney, 2010). We will present a new modelling approach to pollinator preferences in a field of wild flowers without the influence of agriculture.

Materials and Methods

Field Study Data

The data analyzed in this study is from a field study conducted by Vera Pfeiffer in the summer of 2011 and by the students at the Eco-Informatics Summer Institute (EISI) over the summers of the years 2012-2014. The field observers recorded plant and pollinator interactions in 18 meadows in the H.J. Andrews Experimental Forest, involving 109 plants and 293 pollinators over 4 years. During each year, observations were collected at a subset of the 18 total meadows. Each meadow was visited by researchers about 5 times each summer. Each visit is called a *meadow-watch*. For example, Figure 1 shows the interactions between plants and pollinators during one meadow watch.




Figure 1: This figure shows the plant-pollinator interactions recorded at the Lookout meadow, watch 2 in 2014. Each row represents a plant that was available during the meadow-watch and each column represents a pollinator that was observed interacting with a plant during the meadow-watch. The darker a cell is shaded, the more interactions were observed between the respective plant and pollinator. The most abundant interactions were between *Apis mellifera* and *Gilia capitata*, *Apis mellifera* and *Eriophyllum lanatum*, and *Bombus mixtus* and *Lupinus laxiflorus*. There were a total of 1729 interactions recorded at this meadow-watch.

In each meadow-watch, observations of plant-pollinator interactions were made in 10 plots in the meadow during 15 one-minute intervals. The record for each interaction includes the plant, the pollinator, the meadow-watch, the plot, and the number of times the interaction was seen during that minute. At each interval, the temperature, time of day, and cloud cover were recorded. For our purposes, we will use $N_t(i, j)$ for the total number of interactions observed between plant i and pollinator j over all 10 plots and all 15 one-minute intervals of meadow-watch t .

Before every meadow-watch, researchers surveyed the plant species in all 10 plots in the meadow watch, noting the plots where the plants were located as well as how many flowers of that plant species were blooming. This data was compiled into a separate dataset for plant availability. We will use $A_t(i)$ to indicate the total number of flowers of plant species i during meadow-watch t .

Models

Multinomial Model

The goal of the multinomial model is to statistically determine which plants a pollinator prefers from the field study data by fitting a multinomial distribution. In this model, a given pollinator assigns *preference scores* (also referred to as scores and by ϕ) to each plant species indicating how much the pollinator likes that species. A higher score indicates that the plant species is preferred more than a plant species with a lower score. Our goal in this model is to find ϕ_j , which represents the score function for some pollinator j .

In the multinomial model, we can think of the pollinator “rolling a die” with the names of a plant species on each side of the die to determine which plant species it will visit next. However, this die is not evenly weighted. Each side of the die will be weighted differently according to a score that depends on the preference score and the relative abundance of the plant in the meadow. The preference score for each plant is fixed for the full dataset. We then use the interaction data for that pollinator species and the probabilities of the pollinator visiting each plant to calculate the likelihood of the set of observed interactions.

First consider the set $X = \{1 \dots P\}$ which enumerates all plant species. On some meadow watch t , only a subset of X is available. The availability of plants at meadow-watch t can be denoted by a vector A_t of size P where the number of flowers of plant species i is denoted by $A_t(i) \in \mathbb{N}$. Now, we know that pollinator j makes a collection of visits on meadow-watch t , $V_t = \{N_t(1), N_t(2), \dots, N_t(P)\}$. We denote the number of times j visits i by $N_t(i) = \sum_{k=1}^{K_t} I[V_t(k), i]$ where $I[V_t(k), i]$ is an indicator function. The total number of visits made is K_t . The indicator function is equal to 1 when $V_t(k)$ is equal to i and 0 otherwise. In this collection, $v_k \in \{1, \dots, P\}$ and v_k represents the k^{th} plant species that was visited by pollinator j on meadow-watch t .

The pollinator assigns a score for each plant i which we will denote as the score function $\phi_{ij} \in \mathbb{R}$. This function gives a real-valued score to plant i indicating the

pollinator j 's relative preference for i . Then the probability that pollinator j 's k^{th} visit was to plant i is:

$$\theta_i(A_t, \phi_j) = P(v_k = i | A_t, \phi_j) = \frac{A_t(i) \exp(\phi_{ij})}{\sum_{i'=1}^P A_t(i') \exp(\phi_{i'j})}.$$

The likelihood of observing the visits V_t visited by pollinator j on meadow-watch t is:

$$L(V_t; \phi_j) = \frac{K_t!}{N_t(1)! \dots N_t(N)!} \prod_i \theta_i(A_t, \phi_j)^{N_t(i)}.$$

Over the entire summer, the likelihood of the observed data is:

$$L(V_t; \phi_j) = \prod_t \left(\frac{K_t!}{N_t(1)! \dots N_t(N)!} \prod_i \theta_i(A_t, \phi_j)^{N_t(i)} \right).$$

As the resulting likelihood values will be extremely small, we will want to work in log-space. Therefore, we will use the loglikelihood of the observed data:

$$\begin{aligned} LL(V_t; \phi_j) &= \sum_t \left(\log\left(\frac{K_t!}{N_t(1)! \dots N_t(N)!}\right) + \sum_i N_t(i) \log(\theta_i(A_t, \phi_j)) \right) \\ LL(V_t; \phi_j) &= \sum_t \left(\log(K_t!) - (\log(N_t(1)) + \dots + \log(N_t(N))) \right. \\ &\quad \left. + \sum_i N_t(i) \log(\theta_i(A_t, \phi_j)) \right). \end{aligned}$$

We also introduce a regularization term to reduce over fitting the model. The parameter lambda will be calibrated, via cross-validation, to optimize the ability of the model to generalize to new meadow-watches. Our final likelihood model will be:

$$PLL(V_t; \phi_j) = LL(V_t; \phi) - \lambda \sum_i \phi_{ij}^2$$

The ϕ_j parameters we choose will be the ϕ_j which maximizes the $PLL()$ function (penalized log likelihood) for the interactions we observe in the field data.

Originally, we planned to use a multinomial model in which we truncated the interaction counts to 0 and 1 to indicate whether the interaction occurred or not. However, the preliminary work showed that this model cannot be applied to all pollinators in general. Please see Appendix A for more details about this model and the results from this model.

Traits Model

The traits model investigates whether the traits of the available plants can predict the preferences of the pollinators. We take two approaches with this model. The first is a two-phase approach where we first fit the score function and then use the traits to predict the scores. The second is a single combined model to see how well the preferences can be expressed in terms of the traits.

For simplicity, we chose a linear model where the traits of the plant species are represented by the matrix $T = \{1, t_1, \dots, t_n\}$ with one row for each species where 1 represents the intercept and w_0, w_1, \dots, w_n are the corresponding coefficients. The term w_0 is the intercept. Our goal in this model is to find w_0, w_1, \dots, w_n which are constant over all the plants.

We can now express ϕ_{ij} , the score given to plant i by pollinator j , in terms of the traits.

$$\phi_{ij} = w_0 + w_1 t_{1i} + w_2 t_{2i} + \dots + w_n t_{ni}.$$

Two-Phase Approach

We will first consider $\hat{\phi}_{ij}$, the score given to plant i by pollinator j , that we fit to the data using the multinomial model. We will use a linear regression to find the values of the coefficients. The linear regression may provide insight into which, if any, of the floral traits seem to affect the preference scores of the pollinator as expected. If the regression shows that there is a correlation between the plant's traits and the score given to the plant, we can use the combined model explained in the next section.

Combined Model

In this model, we will reparameterize the multinomial model such that the probability is calculated using the traits rather than the ϕ score. In that case, the probability in the multinomial model will be:

$$\theta_i(A_t, \hat{\phi}) = P(v_k = i | A_t, \phi) = \frac{A_t(i) \exp(w_0 + w_1 t_{1i} + \dots + w_n t_{ni})}{\sum_{i'=1}^N A_t(i') \exp(w_0 + w_1 t_{1i'} + \dots + w_n t_{ni'})}.$$

We can then continue to evaluate the multinomial model as a function of the traits of a plant. We will use gradient descent to find the coefficients w_0, w_1, \dots, w_n .

The Traits

We will be evaluating the effect of 11 traits of a flower on the score function. The values for most of the traits are categorical values. R handles these categorical values automatically for us. The 11 traits are the following:

1. Biomass/flower: Estimate of reward per flower (inflorescence) per species. Measured by multiplying the height, width, and length of the flower.
2. Visibility: The color of the flowers. The values for this trait are either “bright” or “not bright”.
3. Closed: Flowers can either be closed or open. Closed flowers can only be pollinated by the pollinators that are skilled at prying the flower open. Open flowers do not require pollinators to have that skill.
4. Pendant: Flowers are either suspended or erect. Erect flowers bloom facing up. Suspended flowers open to the side or upside down and pollinators must have the skill to approach the flower from the bottom.
5. General Tube Shape: The general shape of the plant causes some pollinators to be excluded by the plant. Plants are categorized as severe exclusion, moderate exclusion, and poor exclusion depending on the width, length, and shape of the plant’s tube.
6. Life Form: This trait categorizes plants as annuals or perennials.
7. Flower Form: This trait states if the flower is a bowl or plate shaped flower or if exclusions apply.
8. Pollen Size: The value of this feature is “okay” if the pollen of the plant is of average size and “difficult” if the pollen is either small or large. If the pollen is too small or too large, bees are unable to harvest the pollen.
9. Diel: The value of this feature is “yes night” if the plant is ever open during the night time and “no night” if the plant is only open during the day.
10. Platform: This feature can have a value of “strong”, “weak”, and “not useful”. The value “not useful” indicates that the plant is not useful for mating purposes, because the flower does not provide the visibility (advertising) necessary for an insect to find a mate. The value “strong” indicates that the advertising value for this flower is very high, and the value “weak” indicates that the flower does not have very much advertising value.
11. Feebleness: This feature measures how strong the stem of the plant is. If the value of this feature is “feeble”, it means that large insects will break the flower if they try to land on it. If the value is “strong”, the size of the pollinator may also come into play.

Simulated Data Generation

It is necessary to evaluate the models with simulated data in order to debug and understand them. Before we proceeded with the field data, we wanted some reassurance that the model we are working with was doing a good job in the ideal situation where the assumptions of the model are being met. Therefore, in our initial evaluation of the models, we generated simulated plant-pollinator interactions.

To do this, we sampled the simulated data from the distribution specified by the multinomial model using three score functions that we chose to reflect different kinds of pollinators. We generated simulated data for all the meadow-watches over all 4 years and used the availability for plant i at meadow-watch t ($A_t(i)$) that is reported in the field data.

To generate the data, we specified three different score functions. The “normal” score function, representing a generalist pollinator with some preferences, was based on random numbers generated in the range $[-2, 2]$ with a zero-mean. The “half-special” score function, representing a specialist pollinator that prefers a few plants much more than all the others, had all but 3 plants with a score of -2 . The 3 preferred plants had a score of 2. The “true special” score function, representing a specialist pollinator that prefers one plant above all others, had only one plant with a score of 2. All other plants had a score of -2 .

Using these three score functions and the plant availability provided in the *anthesis* data from the field, we calculated the probability of the pollinator visiting each plant at that meadow watch. Then, to simulate the plant-pollinator interactions, we used the `rmultinom()` function in R to randomly assign a pre-set number of visits that pollinator makes to a certain plant in a certain meadow-watch. We chose to allow the pollinator to make 50 visits during one meadow-watch after analyzing the number of visits *Apis mellifera* makes during one meadow-watch. This generates the $N_t(i)$ we use in the multinomial and traits models.

After conducting experiments using the simulated data, we ran the models on field data from each year separately to see how the pollinators in our study responded to different plant distributions. We also ran our model on the cumulative set of data.

To simulate data for the evaluation of the traits model, we followed a similar approach. In this simulation, we used the real traits for each flower along with the real availability for each flower at each meadow. We set weight vector W , which acts as the coefficient to the traits, to find the resulting score function ϕ using the equation $T \cdot W = \phi$. We then simulated a pollinator with score function ϕ using the method described above. A new

weight vector \widehat{W} is fitted to the simulated data. We can generate the score function associated with \widehat{W} by using the equation $T \cdot \widehat{W} = \widehat{\phi}$.

Bootstrapping for Confidence Intervals

One concern we have when using the field observations is that we do not know the true distribution of the data, and we only have one dataset. If we had more datasets, we could create confidence intervals for the preference scores that are fitted by the model to each dataset. This would show us a range in which to expect the true value of the preference score and would give us confidence that some plant species are really ranked higher than or lower than other plant species.

The idea behind bootstrap is to use the field observations as a population that approximates the true distribution of interactions. We can then create new samples from the field observations by resampling from the field observations without replacement. We created 200 bootstrapped datasets for each year and the aggregate dataset over all years. To create the bootstrapped datasets, we resampled from the original field interactions in each meadow-watch the same number of times the meadow was used in the original dataset. We then fitted the model to the bootstrapped meadow-watch interactions.

Evaluation

Goodness of Fit—Chi-Squared Test

The Chi-Squared goodness of fit test determines whether the distribution of the sample population matches the distribution predicted by the fitted model for the interactions observed for one pollinator j . To determine if we can reject the null hypothesis, we compare the expected and observed frequency counts of each variable using the test statistic

$$\chi^2 = \sum_{t=1}^L \sum_{i=1}^{P_t} \frac{(O_{it} - E_{it})^2}{E_{it}}.$$

Here, P_t is the number of plant species available at meadow-watch t , L is the number of meadow-watches in which both the pollinator was observed, and $O_{it} = N_t(i)$. E_{it} is the number of interactions we expect between the pollinator and plant species i on meadow-watch t . We then fit $\widehat{\phi}$ to the field data and calculate $E_{it} = \theta_i(A_t, \widehat{\phi}) \sum_i^P O_{it}$.

The total number of terms in our χ^2 test statistic is X , which is the number of possible (meadow-watch, plant) pairs (each pair is a cell) where the pollinator had a probability greater than zero for visiting that plant at that meadow-watch. X is slightly less than LP where P is the total number of plant species in the dataset.

For this test, we must penalize one degree of freedom for each meadow the pollinator was present in (L) because we use the number of visits observed in the meadow rather than estimating this quantity independently. As we are also estimating ϕ for each plant, the number of parameters is P , and we must penalize one degree of freedom for each parameter. Because we fit a regularization parameter lambda, we must penalize one degree of freedom for that. Our final calculation for the degrees of freedom (df) is

$$df = X - L - P - 1 .$$

The chi-squared test compares a null hypothesis (H_0) to an alternative hypothesis (H_1). In this case, we will run two tests. The first test will have a null hypothesis as the uniform score function (the preference score for each plant species is the same) and the alternative hypothesis will be that the true distribution is not the uniform distribution. The second test will have the null hypothesis that is equivalent to the learned probability distribution and the alternative hypothesis will be that the true distribution is not equivalent to the learned probability distribution.

A smaller χ^2 indicates a better fit. We will then compare the χ^2 statistic with the χ^2 distribution to determine the p-value of this test. If the p-value is less than 0.05, we will reject the null hypothesis for that test.

Correlations

We used the Pearson's product-moment correlation coefficient (Pearson's r) and Spearman's rank correlation coefficient (Spearman's ρ) to evaluate the scores fitted to simulated data using the multinomial model in comparison to the true scores (comparing ϕ to $\hat{\phi}$).

Pearson's r is a test statistic that measures the linear correlation between two variables X and Y . An r value of +1 denotes full positive correlation, -1 denotes full negative correlation, and 0 denotes no correlation between the variables. The Pearson r is defined as follows:

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}$$

In this statistic x_i and y_i are elements in the variable sets X and Y . \bar{x} and \bar{y} are the means of each set. When comparing the true scores and the estimated scores, if the model fits well, the statistic will have a value close to 1.

Spearman's ρ is defined as the Pearson correlation coefficient for ranked variables in two sets X and Y . Each element in the sets X and (X_i, Y_i) are turned into ranks x_i and y_i that are calculated using the position of the element in ascending order in the set. If there are ties, the elements with ties are given an average rank by adding up all the positions and dividing by the number of elements with the same value. After the ranks are ascertained, we calculate the Spearman's ρ statistic:

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)}.$$

In this statistic, $d_i = x_i - y_i$, which is the difference between ranks. When comparing the true scores and the estimated scores, we expect to have ρ close to 1.

Likelihood Ratio Test

We used the likelihood ratio test to compare the multinomial model and the traits model. This test is similar to the Chi-Squared test in that it can be used to make a decision between two hypotheses. However, the likelihood ratio test is used to see if the likelihood of a model with a given set of parameters θ_0 has a likelihood significantly different to the likelihood of a model with a different set of parameters θ_1 where θ_0 is the null hypothesis and θ_1 is the alternative hypothesis.

The test statistic $\Lambda(x)$ calculates the ratio between the likelihood of the observed data using the parameters θ_0 , denoted by the function $L(\theta_0|x)$, and the likelihood of the observed data given θ_1 , denoted by the function $L(\theta_1|x)$:

$$\Lambda(x) = -2 * \log \left(\frac{L(\theta_0|x)}{L(\theta_1|x)} \right).$$

A large value of $\Lambda(x)$ suggests that the null hypothesis gives a much better fit to the data than the alternative.

Results

Multinomial Model

Using Simulated Data

Our first step was to test the multinomial model on the simulated data. For each of the three “true” score functions (normal, half-special, and true-special), we conducted 10 trials. In each trial, the simulated pollinator had 50 interactions with the plants in each meadow-watch. We conducted these tests using both the unregularized multinomial model and the regularized multinomial model to assess the effect of regularization. In each test, we compared the score function fitted by the model using the simulated interaction data to the true score function using Spearman’s ρ and Pearson’s r correlation statistics.

Through these experiments we find that the multinomial model works well when dealing with a true normal score function both with and without regularization. We see in both Table 1 and Table 2, both Kendall’s τ and Spearman’s ρ correlations are quite high for the normal score function. Higher values for these two statistics indicate that the score functions when ranked match up well. The higher value for Pearson’s r statistic in Table 2 for the normal distribution indicates that the exact values of the score functions match up better with regularization. We can see the effect of regularization on Pearson’s r in Figure 1 and Figure 2. We believe that the most appropriate correlation measure for this task is Pearson’s r . Our objective is to find the most accurate score function for the pollinator. Pearson’s r is the most accurate measure for that. Kendall’s τ and Spearman’s ρ correlations measure how well the ranks (which are different from the scores) of the pollinators match up.

The biggest problem with the unregularized model is the large negative values assigned to plants that are available to the pollinator, but are never visited. These large negative values are a result of not being able to calibrate the score function to the data. To penalize the model if it assigns large negative values, we will use the regularized multinomial model. We used cross-validation to choose a λ for the penalized multinomial model and found that the best $\lambda = 0.25$. This was the best regularization term across datasets. We also see that without regularization, the fitted score function does not match the true score function well for the half-special and true-special score functions. Pearson’s r improves greatly with regularization for the half-special and true-special score functions.

Table 1: Correlation between the true score function and fitted score function without regularization over 10 trials.

| | Normal | Half-Special | True-Special |
|-----------------------|---|---|---|
| Spearman's rho | Mean: 0.889 Deviation: 0.027 Range: 0.854-0.926 | Mean: 0.283 Deviation: 0.0 Range: 0.283-0.283 | Mean: 0.165 Deviation: 0.0 Range: 0.165-0.165 |
| Pearson's r | Mean: 0.531 Deviation: 0.043 Range: 0.462-0.608 | Mean: 0.221 Deviation: 0.022 Range: 0.178-0.266 | Mean: 0.221 Deviation: 0.032 Range: 0.192-0.299 |

Table 2: Correlation between true score function and fitted score function with regularization over 10 trials.

| | Normal | Half-Special | True-Special |
|-----------------------|---|---|---|
| Spearman's rho | Mean: 0.865 Deviation: 0.021 Range: 0.827-0.901 | Mean: 0.283 Deviation: 0.0 Range: 0.283-0.283 | Mean: 0.165 Deviation: 0.0 Range: 0.165-0.165 |
| Pearson's r | Mean: 0.812 Deviation: 0.025 Range: 0.768-0.858 | Mean: 0.664 Deviation: 0.059 Range: 0.597-0.796 | Mean: 0.560 Deviation: 0.062 Range: 0.485-0.659 |

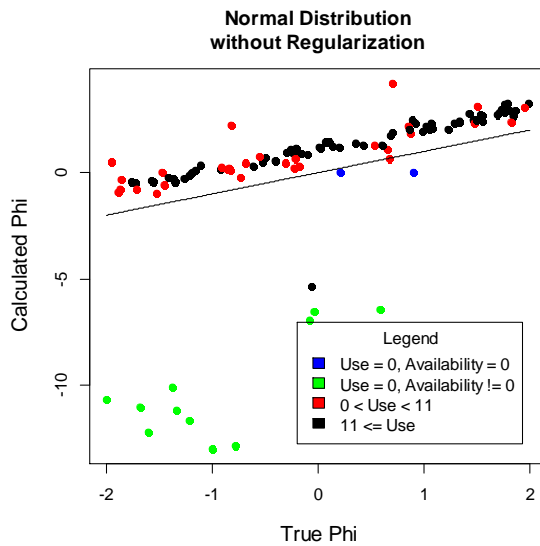


Figure 1: This plot shows the known (true) score function compared to the fitted (calculated) score function for the normal score function. The scores estimated by this model correlate well with the true scores of the pollinator. However, we have a group of large negative values under the $x=y$ line. Also, while the fitted score function generally follows the shape of a linear line, it does not match the $x = y$ line as we would like it to.

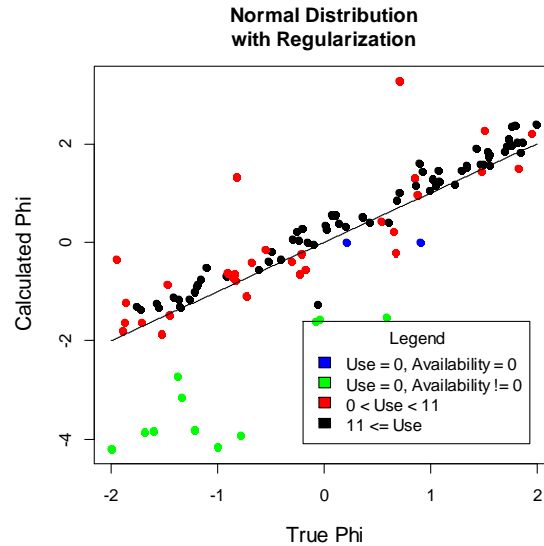


Figure 2: This graph shows the effects of regularization ($\lambda = .25$) on the fitted score function. In comparison to Figure 1, which does not have regularization, the dots on this scatter plot are much closer to the $x = y$ line. The regularization also pulls the large negative values we see in Figure 1 towards their true values which improves the Pearson's r test statistic.

Figures 1 and 2 compare the values of the fitted phi parameters to the true phi parameters in the simulation. A pattern that is interesting to note is that the fitted scores can be separated into 4 categories depending on the availability of the plant and the usage of the plant by the pollinator.

1. Plant is not available, therefore it cannot be used (blue): Because the plant is never available, we cannot learn the pollinator's preferences. Therefore the plant is always given a score of zero. There are a few of these points in the simulated data, but not many. These points are more evident when you have a smaller dataset.
2. Plant is available, but is never used (green): In this case, the pollinator shows a strong preference against that plant species and the plant species is given a very low score. When using maximum likelihood estimation without penalization, the model sends the estimated preference scores to $-\infty$. Regularization forces these values to smaller numbers.
3. Plant is available and used between 1-10 times (red): The points in this category are generally close to the true value of the score for that plant.

4. Plant is available and used more than 10 times (black): This category behaves in the same way as category 3. This category is more useful when analyzing the specialized distributions. The plants that receive the highest scores are usually in this category.

These four categories are easy to see in Figure 1 where there was no regularization. With regularization, these four categories are not as separable. We can also see that the fitted (calculated) score distribution with regularization is much closer to the $x = y$ line. While we have only presented the comparison between the true and fitted score functions for the normal score function, the results are representative for the half-special and true special true score functions. In all subsequent sections, we will use the regularized multinomial model to fit the score functions.

Experiments on the HJ Andrews Data:

To test the model against the field interaction data, we performed goodness of fit tests on the data for each year as well as the aggregate data over all the years. We tested our model on visits made by *Apis mellifera*. The results from the goodness of fit tests are shown in Table 3. All the goodness of fit test resulted in $p = 0$. This indicates that the expected values for the interactions are significantly different from the observed interaction data. We also evaluated the field interaction data against the expected visits for a pollinator using the uniform score function using goodness of fit tests. This test indicated that the field interaction data is significantly different from expected interaction values of a pollinator following the uniform score function.

We also applied the likelihood ratio test to compare the hypothesis that the uniform and fitted preference models are equally good at explaining the observed data versus the alternative that the fitted preference model is better for the aggregate dataset. The test gives very strong evidence against the null hypothesis ($\Lambda = 31483.65$ with 977 df; $p = 0$). This shows that the ratio of likelihoods is extremely unlikely under the null hypothesis that the two models fit the data equally well. This means that the field data is better explained by the multinomial preference model than by the uniform preference model.

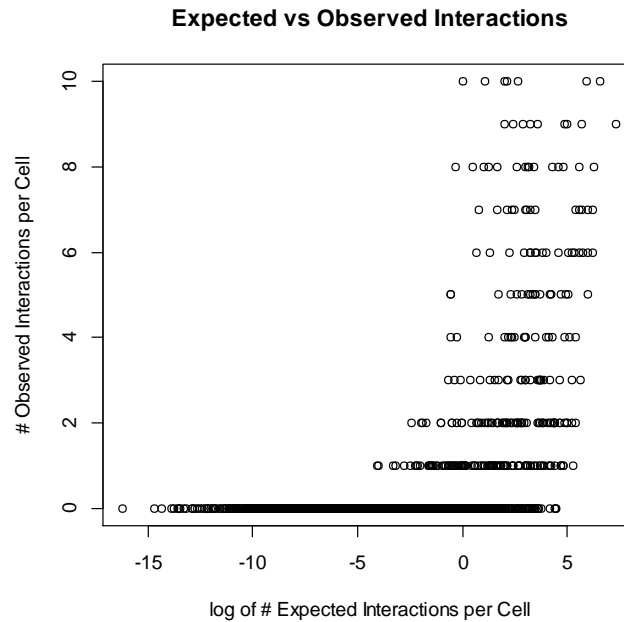


Figure 3: This figure shows the expected number of interactions calculated from the preference scores fitted to the field data using the multinomial model and the observed number of interactions for every possible (meadow-watch, plant species) pair for *Apis mellifera* in the field data. Note that because the expected number of interactions ranged from 10^{-15} to 10^5 , the axis for the expected interactions is in log space. We see here that the expected interactions does not have a strong effect on the observed interactions.

Table 3: The goodness of fit χ^2 received comparing the field interactions of *Apis mellifera* to the expected interactions calculated using either a score function fitted to the field interaction data or the uniform score function. We find that the p-values are extremely low ($p = 0$) in all cases.

| | Goodness of fit for the fitted score function (χ^2) | Goodness of fit for the uniform score function (χ^2) |
|---------------------------|--|---|
| 2011 | 16167407 | 15091.87 |
| 2012 | 1909409 | 11543.29 |
| 2013 | 185040531 | 11103.43 |
| All Years together | 22041400 | 17406.44 |

Table 4 shows the plants that *Apis mellifera* preferred most for each year 2011, 2012, 2013, 2014, and from an aggregation over all four years. We see that *Hydrophyllum occidentale* is the most preferred plant over all of the years. However, it only appears in the top ten most preferred plants of 2011. A closer look at the availability and usage of *Hydrophyllum occidentale* shows that this plant is available only in 2 meadow-watches in 2011 and in 1 meadow-watch in 2012, with a total of 195 flowers across the two years.

However, *Apis mellifera* was only active in one of those three meadow watches (CNE Watch 1, 2011). The only visits *Apis mellifera* made during this meadow watch were to *Hydrophyllum occidentale* although there were many plant species which were more abundant.

Table 4: The top 10 preferred plants for *Apis mellifera* for each year and then for the data aggregated over all of the years. We see that plants that appear in the top 10 most desired plants over all four years are likely to be in the top 10 most desired plants for the aggregate dataset, but not always as with *Gilia capitata*. *Hydrophyllum occidentale* likewise does not fit the pattern. It is in the top 10 for only one year, but also in the top 10 for the aggregate dataset.

| | 2011 | 2012 | 2013 | 2014 | All Years |
|-----------|---------------------------------|-------------------------------|-------------------------------|-------------------------------|---------------------------------|
| 1 | <i>Hydrophyllum occidentale</i> | <i>Rudbeckia occidentalis</i> | <i>Rudbeckia occidentalis</i> | <i>Rosa gymnocarpa</i> | <i>Hydrophyllum occidentale</i> |
| 2 | <i>Symphoricarpos mollis</i> | <i>Luina stricta</i> | <i>Sedum oreganum</i> | <i>Boykinia major</i> | <i>Boykinia major</i> |
| 3 | <i>Cirsium callilepis</i> | <i>Sedum oreganum</i> | <i>Heracleum lanatum</i> | <i>Senecio triangularis</i> | <i>Heracleum lanatum</i> |
| 4 | <i>Rudbeckia occidentalis</i> | <i>Cirsium callilepis</i> | <i>Hypericum perforatum</i> | <i>Senecio integerrimus</i> | <i>Senecio triangularis</i> |
| 5 | <i>Potentilla gracilis</i> | <i>Symphoricarpos mollis</i> | <i>Anaphalis margaritacea</i> | <i>Rudbeckia occidentalis</i> | <i>Senecio integerrimus</i> |
| 6 | <i>Agoseris heterophylla</i> | <i>Potentilla gracilis</i> | <i>Gilia capitata</i> | <i>Hypericum perforatum</i> | <i>Rudbeckia occidentalis</i> |
| 7 | <i>Erigeron foliosus</i> | <i>Phacelia hastate</i> | <i>Perideridia gairdneri</i> | <i>Gilia capitata</i> | <i>Symphoricarpos mollis</i> |
| 8 | <i>Senecio triangularis</i> | <i>Agoseris aurantiaca</i> | <i>Clarkia amoena</i> | <i>Cirsium callilepis</i> | <i>Cirsium callilepis</i> |
| 9 | <i>Gilia capitata</i> | <i>Brodiaea congesta</i> | <i>Erigeron foliosus</i> | <i>Brodiaea congesta</i> | <i>Perideridia gairdneri</i> |
| 10 | <i>Eriophyllum lanatum</i> | <i>Erigeron foliosus</i> | <i>Brodiaea congesta</i> | <i>Potentilla gracilis</i> | <i>Solidago canadensis</i> |

Likewise, we have the top scoring plants for *Bombylius major* and *Eristalis hirtus* for each year and the aggregate dataset in Table 5 and Table 6 respectively. These two pollinators behave quite differently from *Apis mellifera*, but it is interesting to note that these three pollinators all share at least a few plants between their top 10 lists. For example, *Gilia capitata* and *Sedum oreganum* appear in some top 10 list for all three pollinators.

Table 5: The top 10 preferred plants for *Bombylius major* for each year and then for the data aggregated over all of the years. As with *Apis mellifera*, the most preferred plants for *Bombylius major* in the aggregate dataset over all years are all top scoring plants in other years. No new plants are introduced in the top 10 that have not been seen before. However, unlike *Apis mellifera*, *Bombylius major*'s list of preferences changes quite a bit every year. There is only one plant, *Gilia capitata*, that is scored highly over every year.

| | 2011 | 2012 | 2013 | 2014 | All Years |
|-----------|----------------------------------|------------------------------|--------------------------------|--------------------------------|----------------------------------|
| 1 | <i>Gilia capitata</i> | <i>Gilia capitata</i> | <i>Penstemon procerus</i> | <i>Brodiaea congesta</i> | <i>Brodiaea congesta</i> |
| 2 | <i>Mimulus tilingii</i> | <i>Stellaria crispa</i> | <i>Gilia capitata</i> | <i>Fragaria virginiana</i> | <i>Gilia capitata</i> |
| 3 | <i>Arctostaphylos nevadensis</i> | <i>Navarretia divaricata</i> | <i>Montia parvifolia</i> | <i>Galium oreganum</i> | <i>Arctostaphylos nevadensis</i> |
| 4 | <i>Phlox diffusa</i> | <i>Penstemon procerus</i> | <i>Mimulus tilingii</i> | <i>Gilia capitata</i> | <i>Senecio integerrimus</i> |
| 5 | <i>Sedum spathulifolium</i> | <i>Arenaria capillaris</i> | <i>Agoseris heterophylla</i> | <i>Claytonia lanceolata</i> | <i>Claytonia lanceolata</i> |
| 6 | <i>Arenaria capillaris</i> | <i>Vicia americana</i> | <i>Sedum oreganum</i> | <i>Senecio integerrimus</i> | <i>Agoseris heterophylla</i> |
| 7 | <i>Montia parvifolia</i> | <i>Phlox diffusa</i> | <i>Delphinium nuttallianum</i> | <i>Penstemon cardwellii</i> | <i>Mimulus tilingii</i> |
| 8 | <i>Delphinium nuttallianum</i> | <i>Cerastium arvense</i> | <i>Potentilla gracilis</i> | <i>Erigeron foliosus</i> | <i>Phlox diffusa</i> |
| 9 | <i>Calochortus subalpinus</i> | <i>Eriophyllum lanatum</i> | <i>Phlox gracilis</i> | <i>Lupinus laxiflorus</i> | <i>Arenaria capillaris</i> |
| 10 | <i>Penstemon cardwellii</i> | <i>Sedum oreganum</i> | <i>Mimulus nanus</i> | <i>Delphinium nuttallianum</i> | <i>Montia parvifolia</i> |

These top 10 lists also give us some insights about the behavior of these pollinators over the years. *Apis mellifera* seems to be quite stable because the many of the same plants appear over and over again in the top 10 lists over the years. However, *Bombylius major*'s behavior is much less predictable over the years. *Eristalis hirtus*, on the other hand, seems to be quite stable over the years, but then does not behave as expected in the aggregate dataset of all the years combined. In particular, *Perideridia gairdneri* and *Aster ledophyllus* show up in *Eristalis hirtus*'s top 10 list for all years, but neither shows up in the top ten list for a single year.

Table 6: The top 10 preferred plants for *Eristalis hirtus* for each year and then for the data aggregated over all of the years.

| | 2011 | 2012 | 2013 | 2014 | All Years |
|-----------|------------------------------------|-------------------------------|------------------------------|-----------------------------|------------------------------------|
| 1 | <i>Eriophyllum lanatum</i> | <i>Eriogonum nudum</i> | <i>Eriogonum compositum</i> | <i>Cirsium callilepis</i> | <i>Zigadenus venenosus</i> |
| 2 | <i>Polygonum phytolaccaefolium</i> | <i>Eriophyllum lanatum</i> | <i>Erigeron foliosus</i> | <i>Zigadenus venenosus</i> | <i>Cirsium callilepis</i> |
| 3 | <i>Erigeron foliosus</i> | <i>Calochortus subalpinus</i> | <i>Zigadenus venenosus</i> | <i>Eriophyllum lanatum</i> | <i>Polygonum phytolaccaefolium</i> |
| 4 | <i>Zigadenus venenosus</i> | <i>Zigadenus venenosus</i> | <i>Eriogonum umbellatum</i> | <i>Hypericum perforatum</i> | <i>Eriophyllum lanatum</i> |
| 5 | <i>Potentilla glandulosa</i> | <i>Potentilla glandulosa</i> | <i>Potentilla glandulosa</i> | <i>Gilia capitata</i> | <i>Calochortus subalpinus</i> |
| 6 | <i>Aster ledophyllus</i> | <i>Angelica arguta</i> | <i>Perideridia gairdneri</i> | <i>Potentilla gracilis</i> | <i>Sedum oreganum</i> |
| 7 | <i>Calochortus subalpinus</i> | <i>Eriogonum umbellatum</i> | <i>Potentilla gracilis</i> | <i>Phlox gracilis</i> | <i>Potentilla glandulosa</i> |
| 8 | <i>Clarkia amoena</i> | <i>Erigeron foliosus</i> | <i>Sedum oreganum</i> | <i>Rumex acetosella</i> | <i>Erigeron foliosus</i> |
| 9 | <i>Ligusticum grayi</i> | <i>Ligusticum grayi</i> | <i>Eriophyllum lanatum</i> | <i>Ligusticum grayi</i> | <i>Perideridia gairdneri</i> |
| 10 | <i>Gilia capitata</i> | <i>Gilia capitata</i> | <i>Solidago canadensis</i> | <i>Senecio integerrimus</i> | <i>Aster ledophyllus</i> |

Bootstrapping from the HJ Andrews Field Data

A conclusion from our previous testing of the observations from HJ Andrews Forest is that we do not have enough data to have confidence in our results. To explore further, we used bootstrapping to generate 200 new datasets that come from the same distribution as the original datasets. On each new dataset, we calculated the score given to each plant for each year as well as all years and then calculated a 95% confidence interval for the scores of the plants. The following scatterplots (Figures 4 to 8) show the confidence intervals generated for each year.

Figure 4: This plot show the confidence intervals for the preference scores of the plants in 2011 ordered from the plants with the highest upper-bound of the confidence interval to those with the lowest upper-bound. What we notice is that most of the plants have very small confidence intervals. We also see that the confidence intervals of the scores for the top 11 plants do not overlap with the confidence intervals of the plants with lower scores. We also see that the more common plants such as *Rudbeckia occidentalis* and *Rosa gymnocarpa* have larger confidence intervals, as expected




Figure 5: This plot shows the confidence intervals for the preference scores of the plants in 2012 ordered from the plants with the highest upper-bound of the confidence interval to those with the lowest upper-bound.



95% Confidence Intervals using Bootstrapping for 2012

Figure 6: This plot show the confidence intervals for the preference scores of the plants in 2013 ordered from the plants with the highest upper-bound of the confidence interval to those with the lowest upper-bound.


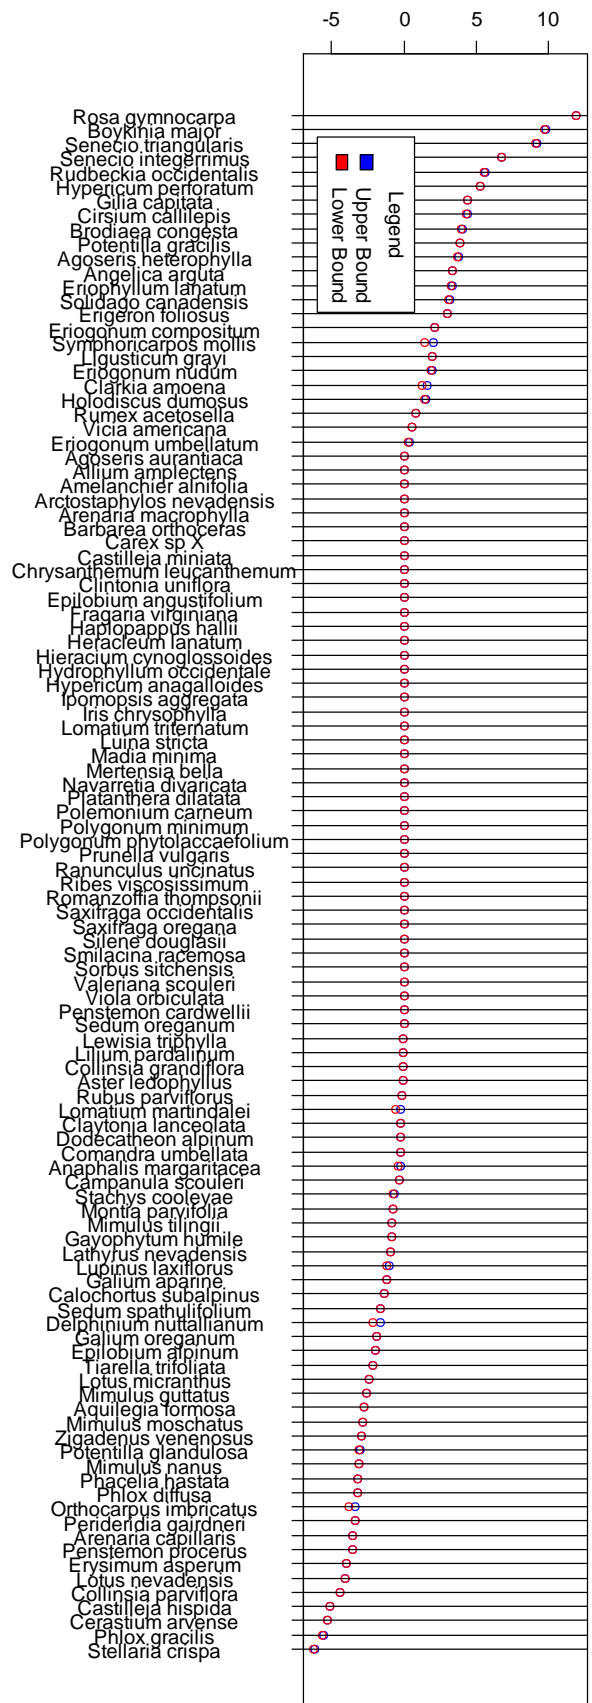



Figure 7: This plot show the confidence intervals for the preference scores of the plants in 2014 ordered from the plants with the highest upper-bound of the confidence interval to those with the lowest upper-bound.



95% Confidence Intervals using Bootstrapping for 2014

Figure 8: This plot show the confidence intervals for the preference scores of the plants in the aggregate dataset over all of the years ordered from the plants with the highest upper-bound of the confidence interval to those with the lowest upper-bound.



95% Confidence Intervals using Bootstrapping for All Years

An interesting outcome of the bootstrapping is that the confidence intervals are very close to the same ranking that we received from the multinomial model in the previous section. Also, for many plants, the confidence interval is very small. This shows that we can be confident about the top ten ranked plants that we display in Table 4, Table 5, and Table 6. We also see that the plants that have confidence intervals with the largest ranges are plants like *Rosa gymnocarpa*, *Rudbeckia occidentalis*, and *Aster ledophyllus*.

Traits Model

Research shows that traits of flowers are important in how a pollinator visits them. For example, *Apis mellifera* is a larger insect and it is unable to hover above flowers. Due to these limitations, we would expect that *Apis mellifera* cannot visit flowers that are extremely small or flowers whose stems are unable to hold its weight. However, there may be many other floral features that may have an effect on pollinator preferences.

Preliminary testing

The first question we explored was if the traits of flowers really had any relationship to the score functions. To explore this, we took the actual usage data for *Apis mellifera* and fitted the score function using the multinomial model. We then fitted the weights for the traits to the score function using linear regression, following the two step approach described in the methods section. Our findings are shown in Table 1. This table shows that the traits that affect *Apis mellifera*'s preferences the most are Flower Form (value: "exclusion exists"), Closed (value: "open access"), and Platform (value: "platform"), as the weights for these traits are the greatest. *Apis mellifera* is a medium-to-large insect and cannot visit very small flowers, which explains why the flower form is important to this pollinator. This pollinator also does not have the ability to hover above flowers, explaining why *Apis mellifera* prefers flowers with a platform. Pollinators also require special skills to open closed flowers, so it makes sense that *Apis mellifera* prefers flowers that are open access.

While the effects of the previous traits are expected, some of the traits (Pendant and Pollen Size) do not seem to have the expected sign. A plant with average pollen size should have a positive weight for *Apis mellifera*, because pollen of average size is the easiest to carry for this pollinator. Likewise, *Apis mellifera* does not have the skill to fly up from underneath a flower, so a suspended flower would be expected to have a negative weight. However, the standard errors for the coefficients of the traits are quite large compared to the coefficients themselves, so the true sign could be different from the sign of the fitted coefficient. The exceptions are the Closed feature which is significant at the 0.05 level and the Flower Form feature which is significant at the 0.1 level.

Table 7: Weights for the traits for *Apis mellifera*. The R^2 value for this test is 0.1696.

| Category | Coefficient | Standard Error |
|--|--------------------|-----------------------|
| Intercept | -0.857089 | 1.322 |
| Biomass/flower | 0.001021 | 0.0024 |
| Closed (open access) | 1.070773 | 1.125 |
| Pendant (suspended) | 0.099545 | 1.002 |
| Visible (not bright) | -0.369354 | 0.766 |
| General Tube Shape (poor exclusion) | -0.216557 | 1.051 |
| General Tube Shape (moderate exclusion) | 0.814345 | 0.811 |
| General Tube Shape (severe exclusion) | -0.638057 | 1.386 |
| Life Form (Perennial) | -0.010474 | 0.687 |
| Flower Form (Exclusion exists) | -1.059014 | 0.799 |
| Diel (flower open at night) | 0.272519 | 0.807 |
| Pollen Size (average) | -0.745493 | 0.605 |
| Platform (platform) | 1.211060 | 0.752 |
| Platform (weak) | -0.182099 | 0.989 |
| Exclusion Feeble (strong) | 0.458899 | 0.881 |

To explore the effect the traits have on the score function in more detail, we also created a few plots visualizing the fitted score function (using the regularized multinomial model) as the function of one trait at a time for *Apis mellifera*. We see in Figure 3 that the biomass of the flower has a slight influence on the score, but not much. Figure 4 shows that the feebleness of the flower has a stronger influence on the score for the plant and verifies the result from the linear regression above. Figure 5 likewise shows that the general tube shape has more effect on the score function than biomass of the flower, but less than the feebleness of the flower. From these plots, we believe that the traits do explain some aspects of the preference scores.

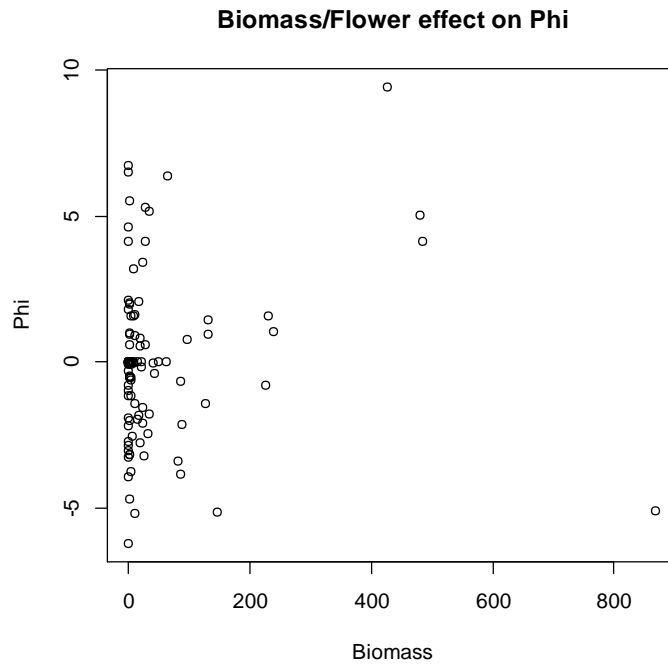


Figure 9: This plot shows flower biomass against the score function. Most of the flowers in this study are quite small and have various scores. The large flower with a very low score at the bottom right of the graph skews the data. Without that point, biomass would have a stronger effect on the score function.

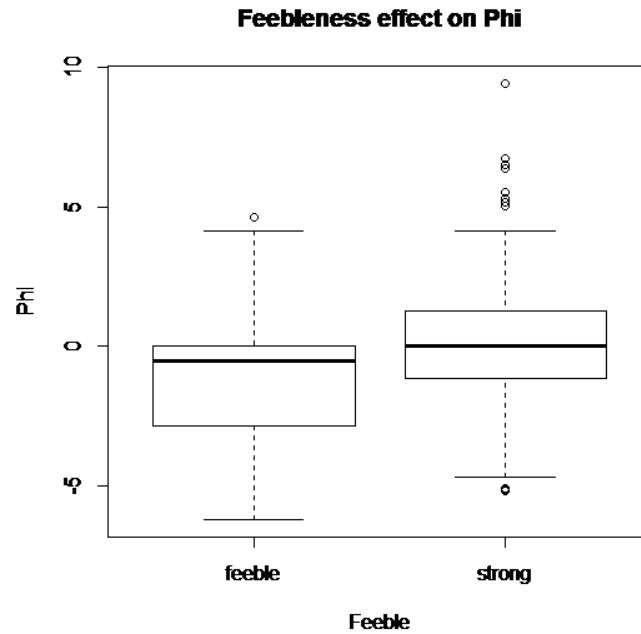


Figure 10: This graph shows how the feebleness of the plant affects the score of the plant. Here we can see that plants with a feebleness value of "strong" correlates with higher scores.

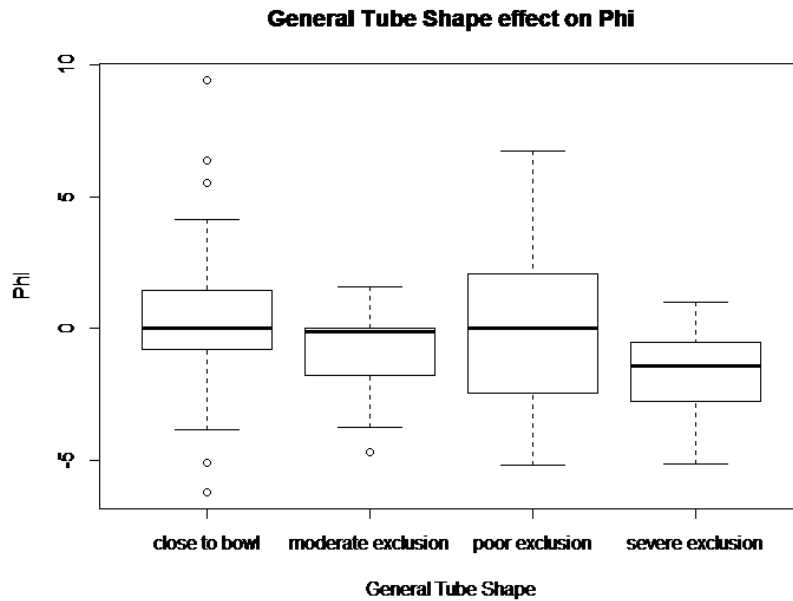


Figure 11: This graph shows the score of the plant compared to the general tube shape of the plant. We see that moderate exclusion and severe exclusion pull the score of the plant down.

Simulated Data Experiments

To make sure the traits model behaves as expected, we generated a weight vector W for the traits model and simulated a pollinator whose preferences ϕ are generated according to that weight vector. We then fit a new weight vector \hat{W} to the simulated data using the single combined model. We can compare W to \hat{W} here and the results of the correlations are in row 1 of Table 8. We also calculated the score functions ϕ and $\hat{\phi}$ using the equation: $T \cdot W = \phi$. We can now compare ϕ to $\hat{\phi}$, and the results can be found in row 2 of Table 8.

Table 8: W is the true weight vector and ϕ is the corresponding true score function which we can calculate from W . We then simulate data using ϕ and fit \widehat{W} to the simulated data. $\widehat{\phi}$ was calculated using \widehat{W} . The following values compare $\widehat{\phi}$ to \widehat{W} and ϕ to $\widehat{\phi}$. The ranges show that there are some cases where W and \widehat{W} correlate very well, but there are other cases when they do not correlate well.

| | Spearman's rho | Pearson's r |
|---|---|---|
| Comparing W to \widehat{W} | Mean: 0.594 Deviation: 0.117 Range: 0.411-0.857 | Mean: 0.811 Deviation: 0.226 Range: 0.330-0.999 |
| Comparing ϕ to $\widehat{\phi}$ | Mean: 0.882 Deviation: 0.085 Range: 0.800-0.992 | Mean: 0.999 Deviation: 0.001 Range: 0.998-1.000 |

Comparison between True and Fitted Weight Vectors




Figure 12: This plot shows the known (true) weights for the traits compared to the fitted (calculated) weights. The weights estimated by this model correlate well with the true trait coefficients for the pollinator. The model generally does very well except for a few points. The model was tested on three weight vectors that were arbitrarily assigned. The color of each point represents the true weight vector that was used. We see that two of the three weight vectors are estimated very well by the model.

We see in Table 8 that when we compare W and \hat{W} , we sometimes get a very good correlation, and sometimes get a correlation that is not as good. Especially with Pearson's r , we see a large range of correlation values. However, when we translate the weight vectors to ϕ and $\hat{\phi}$, we see that the correlations improve. To explore this effect more, we used the field interaction data to test the traits model

Experiments on the HJ Andrews Data

Using the traits model, we report the fitted weight vectors for *Apis mellifera*, *Bombylius major*, and *Eristalis hirtus* in Table 9. Ultimately, these weights do not seem to correlate with what we know about the pollinators themselves. For example, *Apis mellifera* should be attracted to flowers with strong stems (the feebleness trait) and should avoid flowers that are suspended (the pendant trait). Many of the weight values seem like they do not have the correct sign.

Table 9: This figure displays fitted weights for each trait and value combination for *Apis mellifera*, *Bombylius major*, and *Eristalis hirtus* fitted using the traits model on the field interaction data.

| Category | <i>Apis mellifera</i> | <i>Bombylius major</i> | <i>Eristalis hirtus</i> |
|---|-----------------------|------------------------|-------------------------|
| Intercept | -9.881 | 1.301 | -5.738 |
| Biomass/flower | -1.245 | 4.926 | -1.151 |
| Closed (open access) | 3.295 | 1.329 | 7.397 |
| Pendant (suspended) | 2.786 | -2.462 | -3.474 |
| Visible (not bright) | -3.256 | -9.496 | 3.540 |
| General Tube Shape (poor exclusion) | -3.354 | 3.997 | -8.064 |
| General Tube Shape (moderate exclusion) | 9.757 | 3.863 | -1.308 |
| General Tube Shape (severe exclusion) | 6.144 | 3.428 | -1.174 |
| Life Form (Perennial) | 7.188 | 9.902 | -2.628 |
| Flower Form (Exclusion exists) | -7.743 | -2.695 | -6.809 |
| Diel (flower open at night) | 8.001 | -2.269 | 3.518 |
| Pollen Size (average) | -2.103 | -1.802 | -1.237 |
| Platform (platform) | -6.588 | -7.794 | -2.861 |
| Platform (weak) | -2.735 | 1.275 | 9.598 |
| Exclusion Feeble (strong) | -1.500 | -4.978 | -5.211 |

Bootstrapping from the HJ Andrews Dataset

We applied the bootstrapping technique with the traits model with *Apis mellifera* to get a better understanding of the amount of error we have for the traits. We see from Figure 13 that the confidence intervals for the traits are also extremely small. This does not seem intuitive, so before we make any more conclusions on the traits model, more experimentation needs to be done. The confidence intervals may be extremely small because there is not much variation in the bootstrapped datasets. We may have to look back to see if there is a problem with the method we use for the bootstrap.

95% Confidence Intervals using Bootstrapping for All Years




Figure 13: This figure shows the confidence intervals for the weights fitted to the field interaction data for *Apis mellifera* over the aggregate dataset over all years. We created 200 bootstrapped datasets. We see that the confidence intervals for the traits are very small. This can mean one of two things: 1) the weights predicted by the traits model are very accurate or 2) the bootstrapped datasets are not enough to generalize on the true interaction distribution of *Apis mellifera*.

Comparing the Multinomial Model with the Traits Model

To compare the multinomial model and the traits model, we first took one pollinator (we used *Apis mellifera*, *Bombylius major*, and *Eristalis hirtus*) and fitted a score function $\phi^{(1)}$ for the field interactions by the pollinator. We then used the traits model to fit a weight vector W to the same field interactions. These weights are recorded in Table 9. Next we found the score function $\phi^{(2)}$ by using $T \cdot W = \phi^{(2)}$. We then compared $\phi^{(1)}$ and $\phi^{(2)}$ using Kendall's τ , Spearman's ρ , and Pearson's r correlation statistics. We also used the likelihood ratio test to compare the likelihoods of the field interaction data when the pollinator used $\phi^{(1)}$ and when the pollinator used $\phi^{(2)}$. Our results are summarized in Table 10.

In Table 10, we see that $\phi^{(1)}$ and $\phi^{(2)}$ do not correlate well, so the two models are behaving quite differently. However, when we perform a likelihood ratio test on the two resulting score functions $\phi^{(1)}$ and $\phi^{(2)}$, we find that the likelihoods of these score function are not significantly different. The data are not sufficient to reveal any difference in these two models. We cannot draw the stronger conclusion unless we also have some way of measuring the power of the test.

Table 10: Here we fit $\phi^{(1)}$ to the field interactions for the pollinator. We then fit W to the field interactions for the pollinator and calculate $\phi^{(2)}$ using W . We then compare $\phi^{(1)}$ to $\phi^{(2)}$ using Spearman's ρ , Pearson's r , and the likelihood ratio test for each pollinator.

| | <i>Apis mellifera</i> | <i>Bombylius major</i> | <i>Eristalis hirtus</i> |
|-------------------------------------|-----------------------|------------------------|-------------------------|
| Spearman's ρ | 0.209 | 0.110 | 0.276 |
| Pearson's r | 0.207 | 0.119 | 0.389 |
| Likelihood ratio test | -23130.44 | -1636.032 | -1558.958 |
| Δ | | | |

Discussion

Conclusions

In this study, we created two models to simulate and determine the preferences of pollinators: the multinomial model and the traits model. We find that the multinomial preference model performs better than a model that has no preferences and only uses the abundance of the flowers. However, the multinomial model does not explain the behavior of pollinators entirely. The traits model is also not able to account very well for the fitted preferences.

The regularized multinomial preference model explains the visit behavior of several pollinator species much better than a model that assumes that a pollinator has no preferences. When we evaluated the unregularized multinomial model on the simulated data, we found that the model was able to predict only the true score distribution for the “normal” simulation well. However, using regularization improved the multinomial model. Comparing the likelihood of a model without any preferences and the likelihood of the fitted multinomial preference model resulted in a likelihood ratio test statistic $\Lambda = 31483$. Λ is statistically significant and shows the preference model works better than a model without preferences.

While we show that the multinomial model is able to recover the preferences of the pollinator from the simulated data, based on the goodness-of-fit test on the field data, there are still several aspects of pollinator behavior that are not accounted for by the multinomial preference model. When evaluating the multinomial preference model on the field data using a χ^2 goodness of fit test, we find χ^2 is a statistically significant value. This shows that the observed numbers of interactions are significantly different from the expected numbers of interactions. One example of pollinator behavior that is not explained well is the assignment of scores by the multinomial model for rare plant species. For example, when we look at *Apis mellifera*'s preferences over the years to see how the preferences changed, we see that *Hydrophyllum occidentale* appears the top rated plants list only in 2011, but also appears as the top rated plant in the aggregate data set. However, *Hydrophyllum occidentale* is a very rare plant and only appears in the same meadow-watch as *Apis mellifera* once. This lack of information for some interaction pairs could explain why the p-values for the goodness of fit tests on the observed data were so low.

We also wanted to explore how we can predict the observed interactions using the traits of the flowers. This model does not assign scores to each plant. Instead it directly fits the weights on the traits. This steers the multinomial model toward preferences that could be explained by the traits. However, we find that the traits model does not seem to be

finding the true weight vector for *Apis mellifera* and other insects because the values of the weight vector do not seem to have the expected signs let alone the values. Also, a likelihood ratio test between the score function fitted to the multinomial model and the traits model ($\Lambda = -23130$ with 977 df.; $p = 1$) shows us that the traits model and the multinomial model do not behave very differently. We may have to look into other approaches to improve the model even more.

Next Steps

The motivation for this study was to find the plant preferences of a pollinator so that we can study the effects of losing or gaining plant species on plant-pollinator interactions. However, while we have shown that we can take plant-pollinator interaction data and the traits of the plants to find the score function for one pollinator, more work is needed to make the multinomial model predictive and to adapt this model to a plant-pollinator network (with multiple species of pollinators instead of just one).

One idea to make the multinomial model predictive is to incorporate the idea of *burstiness* to explain what the multinomial model is unable to. In the models considered for this study, we do not account for the sequence in which the interactions take place. In essence, we assume that the insect flies up above the flower after each visit and randomly selects a new flower. However, what if once the pollinator finds a flower it gives a high score to, the probability of the pollinator visiting that same plant species is higher? This idea is called *burstiness* and could explain non-random plant-pollinator interactions as well as the effects of competition with other pollinator species. This model is detailed in Appendix B.

Another approach to make the model more predictive is to use the traits data and the interaction data to create a decision tree that classifies the plant as likely to be visited by the pollinator or not likely to be visited by the pollinator. Using this approach, we can find the combinations of traits that attract the pollinator the most. This could also address the problem we have with the current traits model with the weights vector containing un-intuitive values for some of the traits of the flowers.

We would also like to expand the multinomial and traits models so that they can handle multiple pollinator species conducting interactions during the same meadow watch. In the multinomial model, we find the likelihood of the set of visits one pollinator makes. To expand this to multiple pollinator species, we would simply need to multiply the likelihood of the observed visits of all pollinator species. One of the difficulties we might encounter as we expand to multiple pollinator species is keeping the model *density-independent* so that the number of plants or pollinators does not affect the behavior of the model. With the current multinomial and traits models, we account for the density of the

flowers. However, when we incorporate multiple pollinators, we need to account for the density of the pollinator populations as well because in our current dataset, we do not have the ability to count the exact number of pollinators in the meadow during the meadow watch. One idea would be use the number of visits observed as a proxy, but we would not be able to differentiate between pollinators that visit many different flowers and pollinators that are populous. This would affect the density independence of the model.

In this study, we have seen that the model gives different results depending on how much data is observed. The results from an aggregate over all the years are very different from the results over one year. More work should be done on analyzing how these results vary due to the amount of data.

There are also many other ways of dealing with preference. One model that might be interesting to explore in the context of plant pollinator interactions is the Preference-Aversion model introduced by Franco (2013). This model uses a distinction between biological need and a pollinator's desire to learn the preferences of a pollinator in the presence of competition and other non-ideal situations. A model like this could help us account for competition and other challenges a pollinator faces in addition to their desire to visit a particular flower.

Bibliography

- Aizen, M. A., Garibaldi, L. A., Cunningham, S. A. & Klien, A. M., 2009. How much does agriculture depend on pollinators? Lessons from long-term trends in crop production. *Annals of Botany*, pp. 1579-1588.
- Aizen, M. A., Sbatino, M. & Tylanakis, J. M., 2012. Specialization and rarity predict nonrandom loss of interactions from mutualist networks. *Science*, pp. 1486-1489.
- Bascompte, J., Jordano, P., Melian, C. J. & Olesen, J. M., 2003. The nested assembly of plant-animal mutualistic networks. *PNAS*, pp. 9383-9387.
- Biesmeijer, J. C. et al., 2006. Parallel Declines in Pollinators and Insect-Pollinated Plants in Britain and the Netherlands. *Science*, pp. 351-354.
- Bluthgen, N., 2010. Why network analysis is often disconnected from community ecology: A critique and an ecologist's guide. *Basic and Applied Ecology*, pp. 185-195.
- Fieffer, V., 2011-2014. *H.J Andrews Experimental Forest Plant Pollinator Interactions Data*. s.l.:s.n.
- Fitzpatrick, U. et al., 2007. Rarity and decline in bumblebees--A test of causes and correlates in the Irish fauna. *Biological Conservation*, pp. 185-194.
- Gomez, J. M. et al., 2008. Spatial variation in selection on corolla shape in a generalist plant is promoted by the preference patterns of its local pollinators. *Proceedings of the Royal Society*, pp. 2241-2249.
- Hutchinson, R. A., Liu, L.-P. & Dietterich, T. G., 2011. *Incorporating Boosted Regression Trees into Ecological Latent Variable Models*. s.l., Proceedings of the Twenty-fifth Conference on Artificial Intelligence (AAAI).
- Johnson, S. D. & Steiner, K. E., 2000. Generalization versus specialization in plant pollination systems. *Tree*, pp. 140-143.
- MacKenzie, D. I. et al., 2002. Estimating site occupancy rates when detection probabilities are less than one. *Ecology*, 83(8), pp. 2248-2255.
- Memmott, J., Craze, P. G., Waser, N. M. & Price, M. B., 2007. Global warming and the disruption of plant-pollinator interactions. *Ecology Letters*, pp. 710-717.
- Moldenke, A., 2015. *Personal communication* [Interview] (16 April 2015).

Moller, A. & Sorci, G., 1998. Insect preference for symmetrical artificial flowers. *Oecologia*, pp. 37-42.

Rands, S. A. & Whitney, H. M., 2010. Effects of pollinator density-dependent preferences on field margin visitations in the midst of agricultural monocultures: A modelling approach. *Ecological Modelling*, pp. 1310-1316.

Rodriguez, I. et al., 2004. Symmetry is in the eye of the 'beeholder': innate preference for bilateral symmetry in flower-naive bumblebees. *Naturwissenschaften*, pp. 374-377.

Tylianakis, J. M., 2013. The Global Plight of Pollinators. *Science*, pp. 1532-1533.

Appendix A: The Binary Preference Model

This appendix summarizes the preliminary studies done with the binary preference model. This model was of interest to us because of the sampling errors that can occur with field observations of plant-pollinator interactions. The presence or absence of an interaction may be more reliable because rare pollinators are only observed a few times and we cannot determine if the observations we see are due to specialization or due to the probability of observing the interaction. We conducted the experiments with this model on synthetic data so that we would know exactly how the model performed.

Unfortunately, this model seems to work only with generalist pollinators. This model is a special case of the multinomial preference model introduced in the Methods section where the interaction counts can only be 0 or 1.

A Multinomial Preference Model

Multinomial Model

In this model, we are trying to determine which plants a pollinator prefers from the field study data. This model is a statistical approach, and our goal in this model is to find ϕ_j , which is the score function for pollinator j . We assume that all the interactions during a meadow watch are performed by an individual pollinator of species j .

First consider the set $X = \{1 \dots N\}$ which enumerates all plant species. On some meadow watch t , only a subset of X is available. The availability of plants at meadow-watch t can be denoted by a matrix A_t of size $N \times 1$ where the number of flowers of plant species i is denoted by $A_t(i) \in \{0, 1\}$. On this meadow-watch, some pollinator j makes K_t visits to flowers, which we will denote as a collection of visits

$V_t = \{v_{t,1}, v_{t,2}, \dots, v_{t,k}, \dots, v_{t,K}\}$. In this collection, $v_k \in \{1, \dots, N\}$ and v_k represents the k^{th} plant that was visited by pollinator j on meadow-watch t . Please note that the notation v_k is for the ease of explaining this model. The order of the K visits is arbitrary.

The pollinator also has a score for each plant which we will denote as the score function $\phi_j \in \mathbb{R}$. This function gives a real-valued score to plant i indicating the pollinator j 's preference for i . Then the probability that pollinator j 's k^{th} visit was to plant i is:

$$P(v_k = i | A_t, \phi_j) = \frac{A_t(i) \exp(\phi_j(i))}{\sum_{i'=1}^N A_t(i') \exp(\phi_j(i'))}.$$

From here on, we will make the dependence of the function $P(v_k = i | A_t, \phi_j)$ on A and ϕ_j implicit and will refer to this function as $P(v_k = i)$. Now, we know that pollinator j makes a sequence of visits on meadow-watch t , V_t . Then the number of times j visits i is $N_t(i) = \sum_{k=1}^{K_t} I(V_t(k), i)$ where $I(V_t(k), i)$ is the indicator function. The indicator function is equal to 1 when $V_t(k)$ is equal to i and 0 otherwise. The likelihood of this particular collection of plants visited by pollinator j on meadow-watch t is:

$$L(V_t; \phi_j) = \frac{K_t!}{N_t(1)! \dots N_t(N)!} \prod_i (P(v_k = i)^{N_t(i)}).$$

Over the entire summer, the likelihood of the observed data is:

$$L(V_t; \phi_j) = \prod_t \left(\frac{K_t!}{N_t(1)! \dots N_t(N)!} \prod_i (P(v_k(j) = i)^{N_t(i)}) \right).$$

As the resulting likelihood values will be extremely small, we will want to work in log-space. Therefore, we will use the loglikelihood of the observed data:

$$\begin{aligned} LL(V_t; \phi_j) &= \sum_t \left(\log\left(\frac{K_t!}{N_t(1)! \dots N_t(N)!}\right) + \sum_i N_t(i) \log(P(v_t(j) = i)) \right) \\ LL(V_t; \phi_j) &= \sum_t \left(\log(K_t!) - (\log(N_t(1)) + \dots + \log(N_t(N))) \right. \\ &\quad \left. + \sum_i N_t(i) \log(P(v_t(j) = i)) \right) \end{aligned}$$

Binary Model

Because the data in our data set does not tell us the frequency at which a pollinator visits a plant—only that it visits the plant—we must change the model because we only know if the interaction occurred or not. With binary data, the calculated $N_t(i)$ would be either zero or one. Therefore we will call this the usage of a plant i by pollinator j $U_i(i) \in \{0, 1\}$. We can liken $U_i(i, j)$ to a function that returns 1 if pollinator j visits plant i in the sequence V_t and 0 if j does not visit i .

Using this information, we can change the probability function to:

$$P(U_t(i, j) = 1 | A_t, \phi) = \frac{A_t(i) \exp(\phi(i, j))}{\sum_{i'=1}^N A_t(i') \exp(\phi(i', j))}$$

Now we must change our likelihood estimate accordingly. $K_t = U$

$$L(V_t; \phi) = \prod_t \left(\frac{K_t!}{N_t(1)! \dots N_t(N)!} \prod_i (P(v_k(j) == i)^{N_t(i)}) \right)$$

$$L(V_t; \phi) = \prod_t \left(\left(\sum_j U_t(i, j) \right)! \prod_i (P(v_t(j) == i)^{U_t(i, j)}) \right)$$

For the synthetic data analysis, we will use the log likelihood function to avoid precision errors:

$$LL(V_t; \phi) = \sum_t \left(\log \left(\sum_i U_t(i, j) \right)! + \sum_i U_t(i, j) \log(P(v_t(j) == i)) \right)$$

Synthetic Data Analysis

The Algorithm

We tested our model on synthetically generated data. This is the algorithm with which we generated data and analyzed the model:

1. Choose the score function for the pollinator.
2. Calculate the probability a pollinator would visit each plant in the meadow
3. Generate random use data for all the meadows using the probability vector. The generated data was frequency. To do this we used the function `rmultinom()` in R. Here we must also specify the number of visits the pollinator during one sampling unit.
4. Truncate frequency use data to binary use data.
5. Use gradient descent to find the optimal values of the parameters using binary use data. We optimized the log likelihood of the interactions. We utilized the `optim()` function in R to perform the gradient descent.
6. Compare the results with the original score distribution using a rank correlation method.

We hard-coded 4 different score functions:

- Normal: a generalist pollinator with some preference
- Uniform: a generalist pollinator with no preference
- True-special: a specialist that prefers one plant above all others
- Half-special: a specialist that prefers a few plants much more than all others

We required the score functions to calculate the probability that a plant will be visited. We also required the availability of the plants for each meadow watch which we pulled from our existing data.

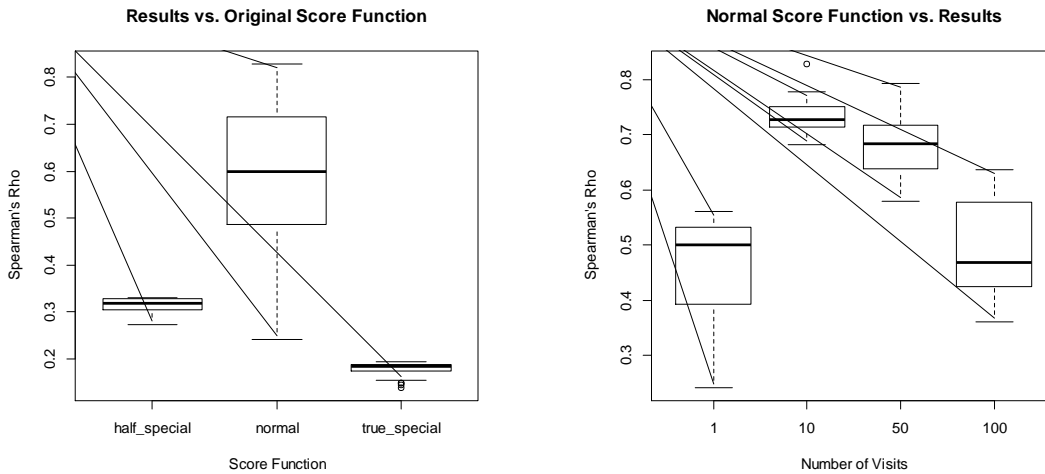
To compare the results from `optim()` and the original score function, we chose to use either Spearman's rho or Kendall's tau rank correlation coefficients. Spearman's rho is better at detecting large discrepancies in the two sets whereas Kendall's tau is more accurate with small sample sizes. Spearman's rho is generally higher than Kendall's tau. For our purposes, I feel that Spearman's rho will be better metric as small discrepancies in ranking are not too important.

We were unable to use the R rank correlation function with our uniform distribution.

The Experiments

We had two variables that we changed: the number of visits the pollinator makes at one time and the score function used by the pollinator. We had 4 options for the number of visits: 1, 10, 50, and 100. We had 3 options for the score function: normal, true special, and half special. We conducted 15 trials for each combination resulting in a total of 180 trials.

Figure 1 shows the results of all the trials for each the original score functions. Both the true special and half special distributions have a very small range and have a low correlation coefficient. On the other hand, the normal score function has a very wide range and also has a very high correlation coefficient at times. To see where this model works the best, I took the data associated with just the normal score function and split it up by the number of visits, resulting in Figure 2. The highest median rank correlation occurs when the number of visits is 10.



Conclusions

This model does not work well with many of different types of score functions. However, it works rather well with a generalist pollinator which has some preferences (i.e. the normal score function). Also, the model seems to work best when the number of visits made by a pollinator at one time is about 10. We believe that when there are not enough visits, we do not have enough information to accurately predict the preference rankings. However, when we have too many visits, rare interactions may be seen, and the truncation from frequency to binary doesn't give us the information about whether a certain interaction is rare or not. This would also drop the correlation with the original rankings.

This model should work well when analyzing species such as *Bombus mixtus* which are known to visit many plants. *Bombus mixtus* has between 10-50 interactions in a meadow for most meadows. However, there are a few meadows there this pollinator has had more interactions. While I was able to calculate the score function for *Bombus mixtus*, I have not yet implemented cross-validation to evaluate the score function.

Appendix B: The Burstiness Model

The Model

The multinomial model and traits model described in the Multinomial Model section treat each bee-flower visit as independent. While this model may explain much about the behavior of insect pollinators according to their preferences, the sequence of visits the pollinators make may be affected by competition from other pollinators as well as the presence of predators.

To incorporate these factors into the model, we use “burstiness”. The idea behind burstiness is to have a transition probability between each visit. Like in the multinomial model and traits model, here we will assume that a single pollinator is making all of the visits for a single meadow watch. Suppose the plant species can be modeled as consisting of a set of C species groups: $\{G_1, \dots, G_C\}$. At each visit, the pollinator has two choices for the next visit:

1. Stay in the same group
2. Leave the group and randomly choose the next group

We will model the pollinator’s behavior as follows:

1. Choose a group g to visit
2. Choose a flower from species $j \in g$ according to the multinomial model (but now *restricted to the flowers belonging to this group* that are available on day t in meadow m)
3. Visit the flower
4. With probability δ go to step 1 and choose a new group. Else go to step 2 and stay with the current group.

The loop in steps 2-3-4 creates “burstiness”.

Now let’s construct the probabilistic model. For step 1, we can apply the multinomial preference model and then compute the probability of each group. Hence, the probability of pollinator i selecting group g is

$$P(g|\phi) = \sum_{j \in g} P(v_k = i | A_t, \phi)$$

Now suppose we had observed the individual visits $V_t = \{v_1, v_2, \dots, v_{K_t}\}$, and let $v = 1, \dots, K$ index those visits. Let g_k denote the group to which species v_k belongs. Then the likelihood of V_t is

$$L(V_t|\phi, \delta) = P(g_1|\phi)P(v_1|g_1, \phi) \prod_{k=2}^{V_t} P(g_k|g_{k-1})P(v_k|g_k, \phi)$$

where the transition probability is defined as follows

$$P(g_v|g_{v-1}) = \begin{cases} (1 - \delta) + \delta P(g_v|\phi) & \text{if } g_v = g_{v-1} \\ \delta P(g_v|\phi) & \text{otherwise.} \end{cases}$$

To simplify the model, we start by assuming that each plant species j belongs to exactly one group g . Now, because each group only has one plant, the probability of choosing a certain plant in a group is always 1. We can remove the terms $P(v_1|g_1, \phi)$ and $P(v_k|g_k, \phi)$ to simplify the likelihood to:

$$L(V_t|\phi, \delta) = P(g_1|\phi) \prod_{k=2}^{V_t} P(g_k|g_{k-1})$$

We can view this as a simulation. The bee flips a coin with probability δ of heads. If it comes up heads, then the bee draws a “new” group according to $P(g_{t,v}|\phi)$. If the coin comes up tails, the bee stays with the current group. Hence, if the state did not change, this could be either because of the $1 - \delta$ probability of tails or because of the probability δ of heads followed by randomly drawing the same group again. If the state does change, it reflects heads followed by drawing a different group.

In practice, we only observe the number of visits $\{N_{t,m}(j)\}_{j=1}^N$, so the likelihood must sum over all unique permutations $\mathcal{P}(V_t)$:

$$L(N_t|\phi, \delta) = \frac{1}{N!} \sum_{\pi \in \mathcal{P}(V_t)} L(\pi(V_t)|\phi, \delta)$$

A Proposed Algorithm

In the formula above, $\pi(V_t)$ is a unique permutation of the original sequence of flower visits. To evaluate this likelihood efficiently, we propose a dynamic program:

change use matrix U to a vector of plant visits V

Note: the ordering of V does not matter here

list = distinct elements of V

lh = calculate the probability of the elements in list.

Note: At this point we have distinct permutations for one visit

for visit in 2:total_use:

list_temp = empty vector


```
lh_temp = empty vector
for entry in 1:length(list):
    v_temp = V - elements present in list[entry]
    for distinct plant in v_temp:
        append [list[entry], plant] to list_temp
        calculate likelihood of this permutation by multiplying transition prob with
        lh[entry] and append to lh_temp
list = list_temp
lh = lh_temp
```

Total_Likelihood = sum(all elements of lh)