# Simulation of Floral Specialization in Bees

**Jessica Oftelie**
Mathematics Department
Iowa State University
Ames, Iowa 50011
joftelie@iastate.edu

May 27, 2004

# Abstract

In this study virtual bees are evolved to visit simulated flowers and gather nectar. The object of this study is to see if floral constancy arises in the virtual bees. Floral constancy, observed in real bees, is the tendency to harvest nectar from only one type of flower. Flowers can only reproduce if they have received the proper pollen type from a visiting bee and so floral constancy is potentially of substantial importance to flowering plants. Two sets of experiments are perfomed. In the first, the virtual bees are evolved for 250 generations, and in the second for 300 generations. The virtual bees are evaluated for floral specialization. Floral specialization is measured as the average of the maximum visits to one flower type divided by the total number of flower visits. The initial hypothesis was that populations with flowers that had nearly equal nectar yields would not specialize, but populations with flowers that had a large difference in yield would specialize in the flower with more nectar. In the first data set, although populations with a choice between nearly equal yields stabilized at non-specialization, the other populations did not behave as expected. Populations that were given one flower with nearly no nectar (less than 0.1000) and a flower with a larger amount of nectar specialized in the flower with nectar, but when given a choice between flowers with any amount of nectar greater than 0.1000, the populations eventually stabilized at non-specialization. In the second data set, populations with one flower type with a nectar level less than .3 were more likely to develop a preference for the other flower, but none specialized.

# 1 Introduction

Bees are an essential part of the flower reproductive cycle for many types of flowers; some flowers cannot reproduce without being visited by a bee. Bees are just as dependent on flowers for their success: the main food resources of bees are nectar (which is stored as honey) and pollen. But how did the relationship between bees and flowers start? Looking back several million years, we can see that there was a time before flowering plants or bees existed, yet today there are a multitude of varieties of both bees and flowers. The fossil evidence currently available indicates that wasps and coniferous plants were around first, followed by bees and flowering plants [5]. Experts hypothesize that bees evolved from wasps who took advantage of pollen available from the plants available at the time, and the plants evolved more efficient ways of luring bees to transfer pollen, but the fossil record is too spotty to say anything for certain about the coevolution of bees and flowers [5].

It is the evolution of bees from wasps, specifically, the evolution of a behavior called floral constancy that is the focus of this study. Floral constancy is the tendency of an individual bee to visit the same type of flower over a period of time. It occurs in polytropic bee species (where the species as a whole visits multiple types of flowers) and varies not only between individuals but also possibly over time in a given individual bee. This paper is the first part of a longer study on the simulation of floral constancy in virtual bee populations.

The foraging behavior of bees seems to be both inherited and learned, in that bees are able to forage on their initial flight, but can improve their efficiency through experience with particular flowers. Bees are able to remember the time and location at which a particular resource (flower species) was found for up to twenty-four hours, and some of the more social bees (honey bees) can communicate that information to their hivemates.

Foraging behavior of a species of bees is classified as monotropic, or visiting only one type of flower, oligotropic, visiting several species of flower, or polytropic, visiting many species of flowers. The tendency for individual bees to visit the same type of flower over a period of time is called floral constancy. According to Roubik, bees forage wherever they find rewarding nectar sources, and the amount of consistence expressed by an individual will depend on the relative abundance of different resources during their life span and (if genetic variations affect floral visitation patterns) over generations [9]. Although at any given time a species of polytropic bees in a particular area, while foraging from a wide variety of flowers, may show a tendency to favor some flowers over others, the preferences of individual members of the species are independent [5]. Observations made by Heinrich in a field of red clover, fall dandelion, and wild carrot showed that while all species of bumblebees foraged from all resources, each species had an apparent preference (measured by the percentage of individuals on a type of plant). When individuals were followed, however, he found that the not all members of a species shared that preference.

Bees have several instinctive behaviors which combine to help them utilize resources at close to optimal levels. The better the quality and quantity of the nectar in the flower they are currently in, the shorter the distance they will fly when they leave. Conversely, after visiting a poor producer, they will fly longer distances [9]. Roubik also states less nectar is removed from flowers in areas that have been determined to be locally plentiful.

While gathering nectar or pollen, bees turn randomly in a circle. They do not recall the direction they came from [10]. Zimmerman states that the change in direction in foraging movement is random for bumblebees that are successful in gathering resources, and Pyke's research finds that the changes occur more often when a profitable patch is encountered.

These two behaviors together make bees more likely to move on after a poor flower and stay in the same general area of a good flower.

A set of experiments conducted by Heinrich illustrate the utility of these behaviors. He covered patches of white clover with screening to allow nectar to accumulate while the surrounding clover was depleted. When the screening was removed, bees in the nectar-rich area probed more florets, stayed for a longer period of time, and made shorter flights with more sharp angular movements than bees in the depleted areas. This had the effect of concentrating their foraging in the areas most likely to be rewarding.

Bees are born with only general nectar-gathering skills, but some flowers require a specialized method to reach the main nectar deposits. This complexity can limit the number of resource types sampled successfully by a bee [6]. The manipulations for any specific flower must be learned through trial-and-error [9]. Inexperienced bees who have moderate success with these flowers on the first try may return to that species and hone that skill, while bees who have little or no success are most likely to ignore that flower in the future in favor of flowers they find easier. This minimizes the amount of time a bee will spend on a resource it cannot take advantage of fully.

When a flower has good nectar production, it is likely that other flowers of the same species also will be good to visit. By learning to recognize one flower species and searching for it first, bees can save time by not visiting other, less productive species [5]. Experiments conducted by Waser on *Bombus* and *Diadasia* found that foragers discriminated among flowers differing solely in color. Foragers of both types tended to visit the same color flower as the one they had just left, while only occasionally sampling the other type. As shown by Waser's experiment, floral constancy is not total for polytropic bee species. Individual bees, while having a major specialization, will from time to time visit other types of flowers (their "‘minor"’ specialization). If they chance upon a flower that is better than their current major, or if their major disappears, they can quickly change specializations. This allows polytropic bees to adapt to some changes in their environment more readily than specialist bees.

Research into the behavior of individual bumblebees has shown that individuals tend to prefer one type of flower over the other available resources, with only occasional visits to other flower types. These individual preferences may be contrary to general species preferences. This is called major and minor specialization [5]. Bees who have chosen a major specialty often follow predictable flight paths, called trap-lines, to the various patches of their chosen flower. Within each patch their path will be random [5]. Since the trap-lines seem to be learned from visible landmarks, the establishment of a predictable trap-line depends on age and a non-uniform landscape.

The many advantages of floral constancy for bees can be viewed as a means of economizing flight and foraging time (and hence energy expenditure) by preferentially utilizing known resources [3]. Many behaviors related to nectar-gathering must be learned by each individual bee. While it is possible for honeybees to communicate the location of a good resource and help with recognition by transmitting scent information, both the initial discovery of the resource and the actual gathering of the nectar are individual pursuits.

# 2 Experimental Design

There area many independent factors which could contribute to the development of floral constancy in bees. The number and distribution of nectar-producing flower species, quality and quantity of the nectar, availability of landmarks, the life span of an individual bee, and the ability of the bee to learn new behaviors are a few possibilities. For this initial study, the virtual bees were given only information about nectar levels and could only decide to drink or not drink from a flower. The landscape was a uniform field of two randomly distributed species of flowers, so there are no landmarks available. Each virtual bee is a representative agent for a hive.

This study was intended to model eusocial bees, where the only reproductive female is the queen. Since the queen mates only one time and can live for several years, the workers from any queen will be siblings. Thus the number of generations models the life span of the queen. The bees are represented by finite state machines that evolve via the evolutionary algorithm described subsequently. In this model, evolution represents learning, not actual biological evolution. The bees are nectar-seeking, but pick up pollen as a secondary resource. They move between floral loci in a random direction, which simulates the rotating movement of bees while feeding in the field. In the initial populations, the bees could only move one square in any direction, but in the second group, the distance moved was inversely proportional to the nectar found.

The flowers used in this study are arranged on a 60x60 grid. For this initial set of runs, two flower types are used which differ only in the quantity of their nectar and have incompatible pollen types. Flowers are initially seeded randomly on the floral grid. Each loci of the landscape can contain one flower or be empty. Empty squares can be seeded by adjacent squares during the flower reproduction cycle, which occurs every "'year"'. For the initial set of populations, nectar quantity was assigned randomly at the beginning of each new population, while the second set assigned the same nectar levels to fifteen different populations. Nectar levels remained constant throughout the evolutionary cycle. Flowers reproduce only if they receive the correct pollen type from a visiting bee. Flowers that do not receive the correct pollen type die off leaving an empty locus.

## Definitions

The virtual bees are represented as finite state automata (FSAs), also called finite state machines (FSMs). A finite state automaton consists of input and output alphabets, a set of states, a transition function that defines how the automaton moves from input and current state to a new state, and a response function that takes the input and current state and produces an output [1]. The finite state automata in this experiment accept flower types as an input and generates a response that is either "drink" or "do not drink". The machines used in this study have four states. An example of a finite state machine of the type used to control the bees is shown in Figure 1.

Evolutionary computation uses the theory of evolution as an algorithm to accomplish a task, usually either one involving the evolution of a feature or the learning of a task [1]. The driving forces of both evolution and evolutionary computation are variation and selection. It uses variation operators known as mutation and crossover on data structures to make random changes in these data structures and blend together parts of different structures. Any algorithm that favors structures with a higher fitness score can be used as a selection

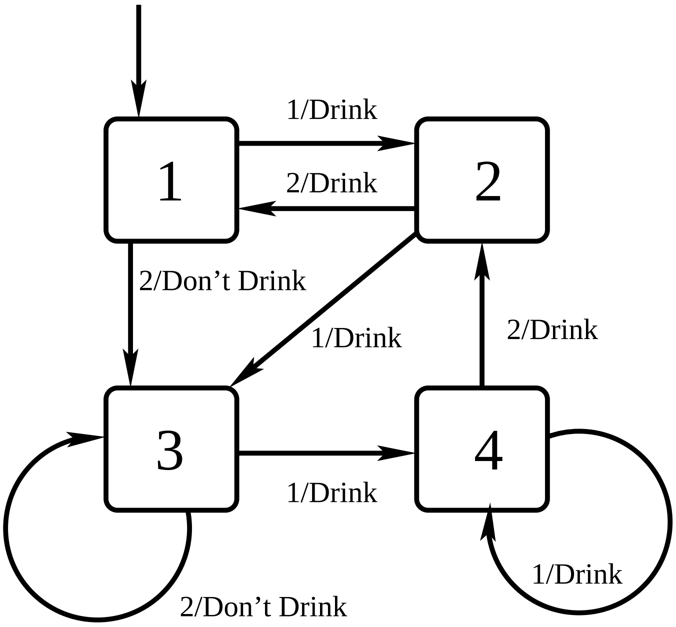

Figure 1: An example of the finite state representation used for virtual bees in this study.

operator. The specific algorithm used in this study is described in Figure 2.

## The Evolutionary Algorithm

The evolutionary algorithm used in this study operates on a population of 60 virtual bees. These bees are thought of as representative agent types for an entire hive. During fitness evaluation the direction of the bee's motion is chosen at random, and the distance is either (for the initial run) exactly one square, or (for the second run) inversely proportional to the nectar gathered. When a bee encounters a flower in the course of its random movement the finite state machines makes a drink/don't drink decision. The bee gains fitness each time they drink from a flower according to the amount of nectar assigned to the floral type, though a given flower can yield nectar only once. In different collections of evolutionary runs different amounts of nectar were used. The amount of nectar contributed by each flower type was generated uniformly at random in the range [0,1] for the initial data set, and chosen from within that range based on the initial results for the second data set. The fitness evaluation of each virtual bee continues for six years. At the end of each year, flowers that were pollinated drop seeds and these seeds sprout to supply the flowers used in the next year.

A virtual bee has 100 pollen "slots" which can be filled with the type of pollen associated with either flower or can be empty. Each time a bee visits a flower three slots are picked. Any pollen in these slots is dropped on the flower and pollen from the flower is placed in these slots. This is a form of proportional pollen selection that makes floral constancy good for flowers but not absolutely necessary for the bees.

During each year the bees floral constancy is computed by the formula

$$Constancy(F_1, F_2) = \frac{Max(F_1, F_2)}{F_1 + F_2} \tag{1}$$

where $F_1$ and $F_2$ are the number of visits to each type of flower. The bee's floral constancy for the six years is the average of its constancies over the individual years. A floral constancy

```
create random population of FSAs
initialize nectar levels for each flower type
     repeat for desired number of generations
     start with a random field of flowers
         repeat for life span of queen
               bees visit flowers and gather nectar
               seeds drop randomly around original flower and sprout
         evaluate fitness as the total amount of nectar collected
         breed if time remains as follows, leaving the original
                     two automata and two children:
               within random groups of 4, choose the two best to be parents
                     and copy them as children
               perform a two-point-crossover in the two children
               perform a random number of mutations on each of the children
```

Figure 2: Evolutionary algorithm for modeling bee learning

of 0.5 is represents a complete lack of floral constancy while a value of 1.0 represents the maximum level possible.

Breeding of the bee automata is performed using single tournament selection with tournament size four. The population is shuffled at random into groups of four bees. In these groups the two virtual bees that gathered the most nectar are copied over the two that gathered the least nectar. The string of states of the newly copied finite state machines undergo two point crossover. After crossover the finite state machines are subjected to 1-3 mutations with the number of mutations chosen uniformly at random. A point mutation changes the initial state of the machine (10%) a drink/don't drink response (50%), or a transition in the finite state machine (40%).

The evolutionary algorithm is in no way meant to model the actual selection pressures on bees or mimic the way bees breed in nature. This study attempts only to see if floral constancy, as observed in nature, will arise in a population of virtual bees being optimized for nectar gathering in the presence of a simple model of pollen limited floral reproduction.

# 3 Results

## First Data Set

The initial hypothesis of this study was that the bees would preferentially visit the type of flowers with more nectar. It was expected that specialization would occur more often when the difference between nectar levels was high. The results of the runs show that specialization does occur in the given situation but only for extreme differences. Stable partial specialization (i.e., a 70% preference for one flower) was not observed. A population not changing from generation to generation either stayed near a floral constancy of 0.5 or near 1.0. A small number of runs exhibited an unstable floral constancy.

| > | A | B | C | D | E |
|---|---|---|---|---|---|
| 0.0 | 0 | 0 | 0 | 18 | 0 |
| 0.1 | 0 | 0 | 1 | 23 | 0 |
| 0.2 | 0 | 0 | 1 | 13 | 0 |
| 0.3 | 1 | 0 | 1 | 13 | 0 |
| 0.4 | 2 | 2 | 0 | 9 | 1 |
| 0.5 | 2 | 1 | 2 | 4 | 0 |
| 0.6 | 1 | 2 | 0 | 2 | 1 |
| 0.7 | 1 | 1 | 1 | 2 | 0 |
| 0.8 | 0 | 0 | 0 | 0 | 0 |
| 0.9 | 3 | 1 | 0 | 0 | 0 |

Table 1: This table shows the number of simulations out of 109 that fell into each of five categories. The rows of the table are indexed by the **difference between the nectar levels** in the two floral types, rounded up to the nearest tenth. The categories are: A - Complete floral specialization occurred. B - High but not complete floral specialization. C - High specialization occurs and then dies away. D - No floral specialization. E - Unstable levels of floral specialization throughout evolution.

A total of 109 evolutionary runs were performed for this data set. The populations evolved can be separated into two types: those that specialized and those that did not. A plot of the floral constancy in one example population of each of these two types is given in Figure 3. The floral seed sprouting numbers for an example of a population of virtual bees that specialized completely in single flower type and a population of virtual bees that did not so specialize are shown in Figure 4. Out of 109 populations, ten specialized in one flower type within 250 generations, 84 reached a stable state around 0.5, and six came close to specializing (had specializations above 0.95) before returning to 0.5. Two populations never stabilized and seven improved steadily but never got above 0.95 consistently. The populations which specialized all had one flower with less than 0.100 ”‘nectar units”’ and one with more than 0.400 ”‘nectar units”’. The closer to 0.1 the lower number was, the higher the other flower's nectar needed to be to specialize. All but three of the 84 populations that stabilized near 0.5 had nectar levels over 0.1. Two of the outliers had the second flower's nectar below 0.3, and one had the lower number close to 0.1 with a higher number close to 0.5. All the populations which approached specialization before returning to near 0.5 had all nectar levels greater than 0.1, and all populations which steadily approached specialization had numbers which resembled the specialized populations.

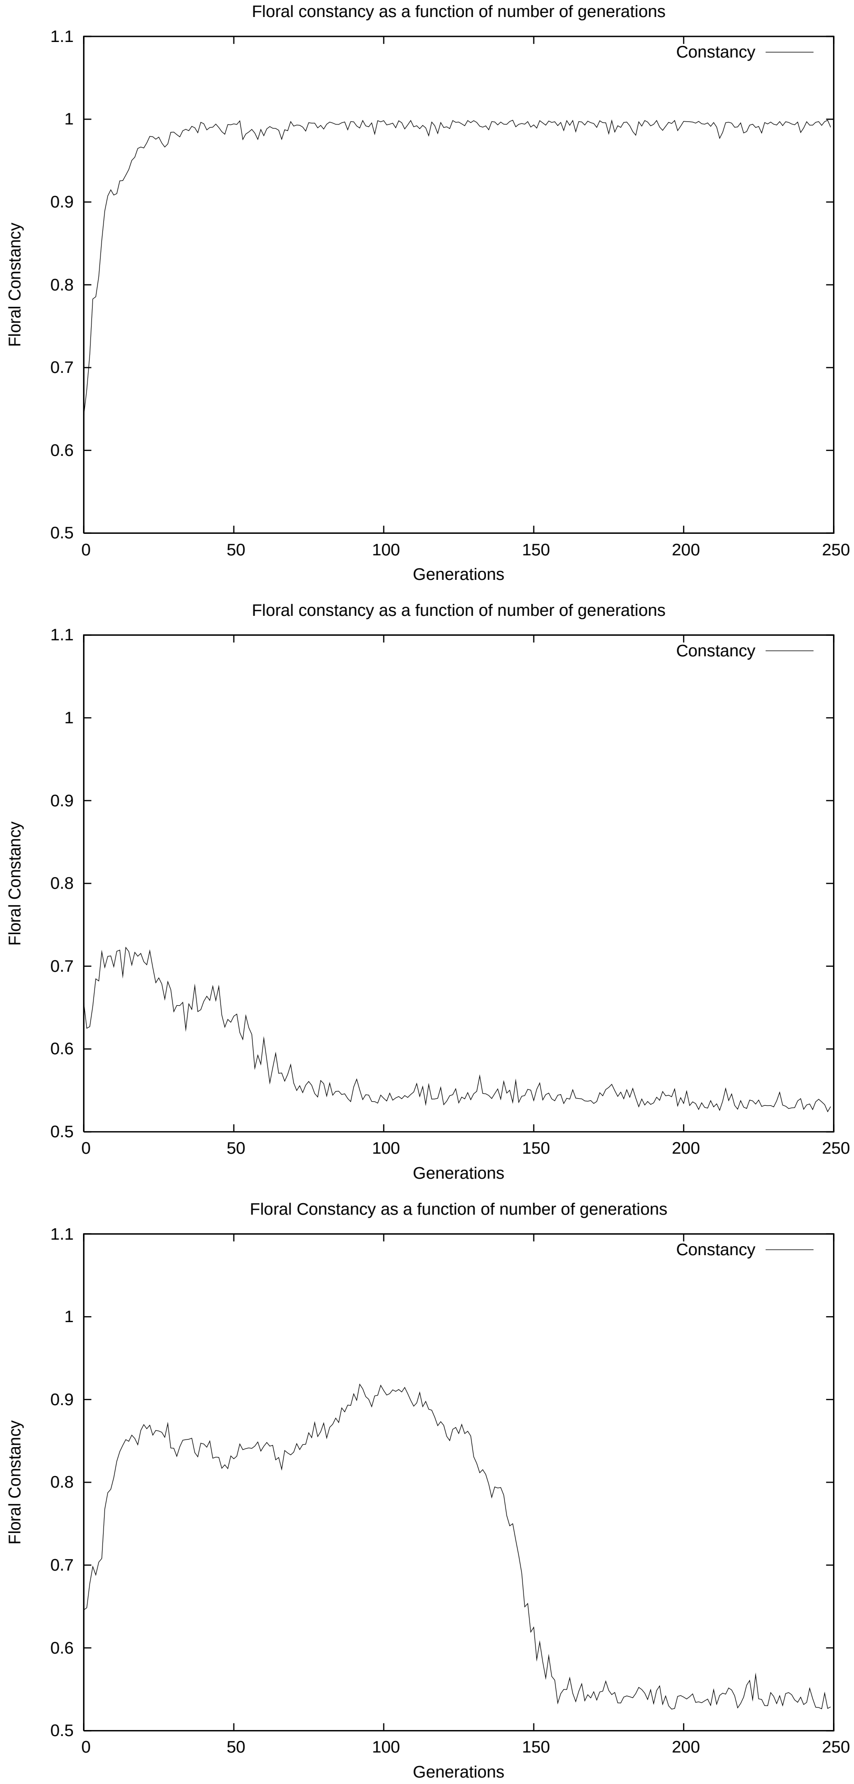

Figure 3: Floral constancy in each generation during the evolution of three population of virtual bees. The top population specializes completely. The middle population fails to specialize at all with the constancy dropping to near 0.5. The bottom population is an example of near specialization relaxing to nonspecialization.

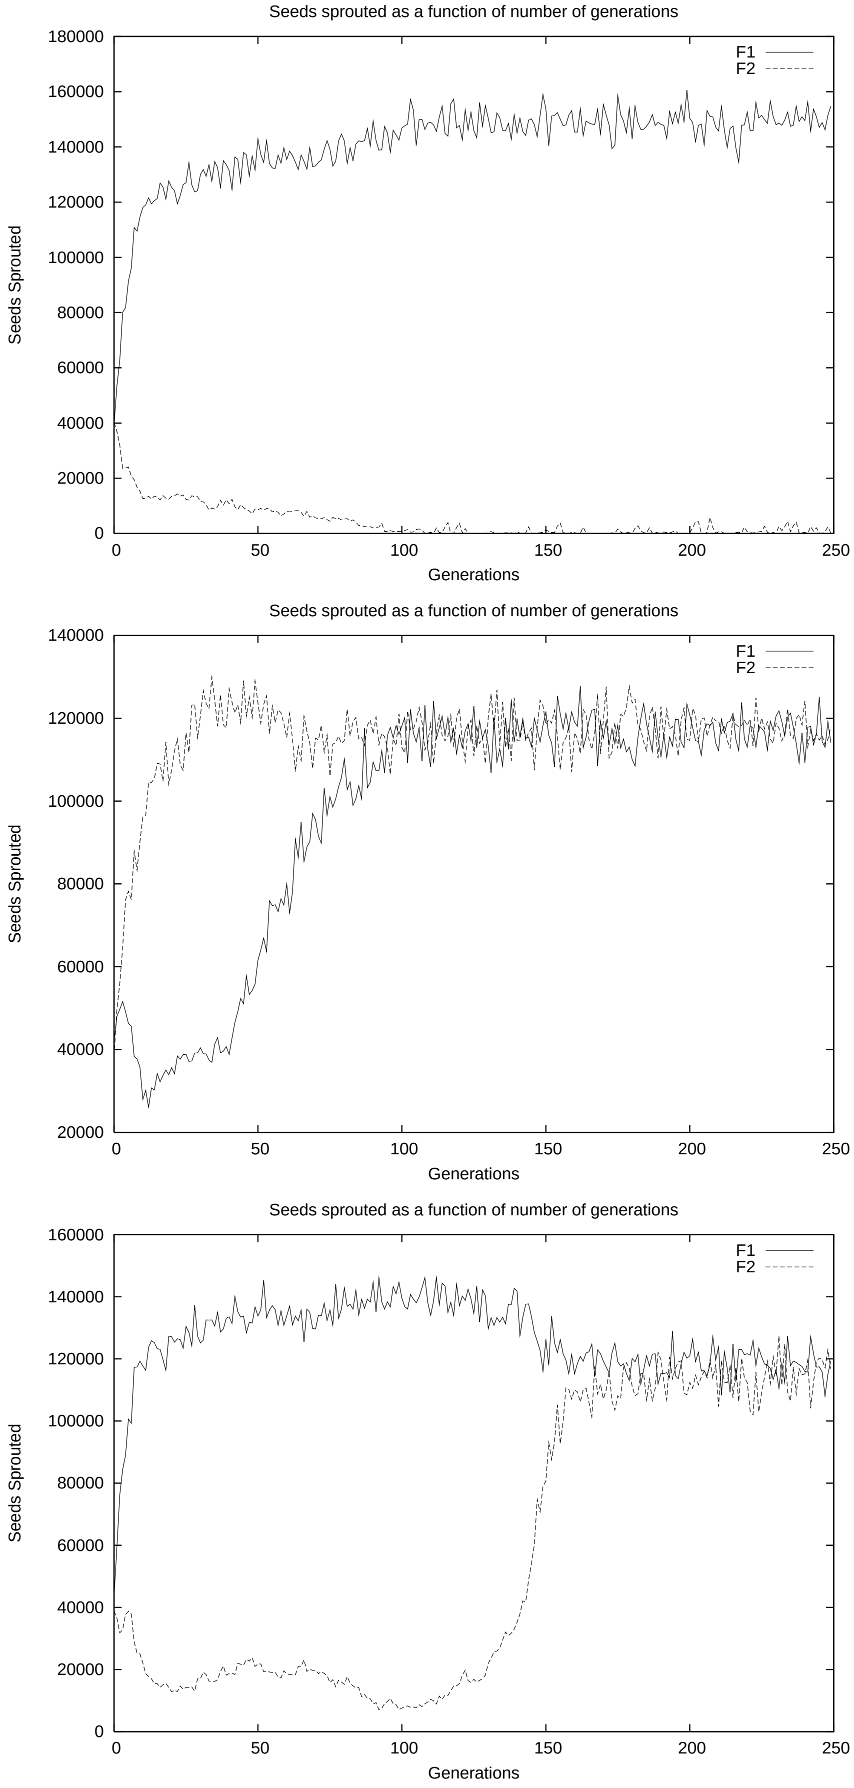

Figure 4: Floral seed sprouting for each of two flower types in each generation during the evolution of a population of virtual bees. The top (specializing) population permits one flower type to die off completely while the center (non-specializing) population sustains both indefinitely. The bottom (eventually non-specializing) population almost lets one flower type die before balancing flower population levels.

| > | A | B | C | D | E |
|---|---|---|---|---|---|
| 0.0 | 10 | 7 | 0 | 3 | 1 |
| 0.1 | 0 | 0 | 2 | 7 | 1 |
| 0.2 | 0 | 0 | 3 | 17 | 0 |
| 0.3 | 0 | 0 | 0 | 20 | 0 |
| 0.4 | 0 | 0 | 0 | 9 | 0 |
| 0.5 | 0 | 0 | 0 | 10 | 0 |
| 0.6 | 0 | 0 | 1 | 8 | 0 |
| 0.7 | 0 | 0 | 0 | 6 | 0 |
| 0.8 | 0 | 0 | 0 | 3 | 0 |
| 0.9 | 0 | 0 | 0 | 1 | 0 |

Table 2: This table shows the number of simulations out of 109 that fell into each of five categories. The rows of the table are indexed by the **lowest availiable nectar** in either of the two floral types, rounded up to the nearest tenth. The categories are: A - Complete floral specialization occurred. B - High but not complete floral specialization. C - High specialization occurs and then dies away. D - No floral specialization. E - Unstable levels of floral specialization throughout evolution.

An unanticipated correlation appeared between levels of seed sprouting and the level of specialization in a population. Sprouting levels of the flower that a population specialized in rise quickly in the same pattern as the rise in specialization, while flowers from populations that do not specialize show sprouting levels that approach near-equal amounts at the same rate as the populations' specialization. This result could be interpreted as the effect of the bees' behavior on the flower population. If this is correct, further studies using more complicated floral models could use these numbers as a measure of floral fitness.

## Second Data Set

In order to study floral constancy in individual bees, a population of polytropic virtual bees is needed. While many populations with no preference were developed, the more interesting case where the population had a non-exclusive preference for one flower type was not observed. Additionally, the populations which changed their behavior from specializing in one flower to being generalists appeared to do so without any pattern or motivation from changes in availability of nectar. The studies by Roubik, Zimmerman, and Pyke cited in the introduction indicate that there are two behaviors involved in the successful gathering of resources – random changes in direction and a flight distance inversely proportional to the quality and quantity of nectar gathered from the current flower. Since only the first criteria was incorporated into the initial data set, the addition of a variable flight distance was the obvious next step in the incremental development of a model of bee behavior.

In the results of the first data set, the populations that encountered large differences in nectar levels were the most difficult to predict, but also were the least represented. To correct this imbalance and to observe the behaviors of different populations in response to the same nectar levels, fifteen populations were run at each of the following nectar level pairs: (0.05, 0.95), (0.1, 0.9), (0.15, 0.85), (0.2, 0.8), (0.25, 0.75), (0.3, 0.7), (0.35, 0.65), (0.4, 0.6), (0.45, 0.55), (0.5, 0.5). Results for the second data set differed more from the first than expected. Populations with low differences in nectar levels behaved much like those in the first set, but populations with large differences did not. No population in the second data set reached full specialization, and none with a difference greater than 0.5 exhibited no preference. Out of the populations with a larger difference, although four managed to reach a 90% preference, most were between 80 and 90%.

| > | A | B | C | D | E |
|---|---|---|---|---|---|
| 0.0 | 0 | 0 | 0 | 0 | 15 |
| 0.1 | 0 | 0 | 0 | 0 | 15 |
| 0.2 | 0 | 0 | 0 | 0 | 15 |
| 0.3 | 0 | 0 | 0 | 0 | 15 |
| 0.4 | 0 | 0 | 0 | 15 | 0 |
| 0.5 | 0 | 4 | 1 | 10 | 0 |
| 0.6 | 3 | 7 | 4 | 1 | 0 |
| 0.7 | 1 | 14 | 0 | 0 | 0 |
| 0.8 | 0 | 15 | 0 | 0 | 0 |
| 0.9 | 0 | 15 | 0 | 0 | 0 |

Table 3: This table shows the number of simulations out of 150 that fell into each of five categories. The rows of the table are indexed by the **difference between nectar levels** in the two floral types, rounded down to the nearest tenth. The categories are: A - > 90% preference for one flower. B - > 80% preference for one flower. C - > 70% preference for one flower. D - > 60% (low) preference for one flower. E - > 50%, (virtually non-existant) preference for one flower.

| > | A | B | C | D | E |
|---|---|---|---|---|---|
| 0.0 | 0 | 15 | 0 | 0 | 0 |
| 0.1 | 1 | 29 | 0 | 0 | 0 |
| 0.2 | 3 | 11 | 5 | 11 | 0 |
| 0.3 | 0 | 0 | 0 | 15 | 15 |
| 0.4 | 0 | 0 | 0 | 0 | 30 |
| 0.5 | 0 | 0 | 0 | 0 | 15 |

Table 4: This table shows the number of simulations out of 150 that fell into each of five categories. The rows of the table are indexed by the **lowest available nectar** in either of the two floral types, rounded down to the nearest tenth. The categories are: A - > 90% preference for one flower. B - > 80% preference for one flower. C - > 70% preference for one flower. D - > 60% (low) preference for one flower. E - > 50%, (virtually non-existant) preference for one flower.

When given the same nectar levels, most of the populations treated the flower types in the same way. The exception was in the populations given flowers with mid-range (from 0.5 to 0.7) differences in nectar levels. Those populations split unevenly in their behavior. There is, however, a corresponding change in the number of seeds sprouted for each flower type. In fact, relative sprouting levels appear to have a better correspondence to levels of preference than nectar levels. The breakdown is as follows:

| Flower A | Flower B | Preference |
|---|---|---:|
| $> 250000$ | $< 30000$ | 90% |
| $> 250000$ | near 50000 | 80% |
| $< 250000$ | $> 50000$ | 70% |
| near 200000 | near 100000 | 60% |
| $< 200000$ | $> 125000$ | 50% |

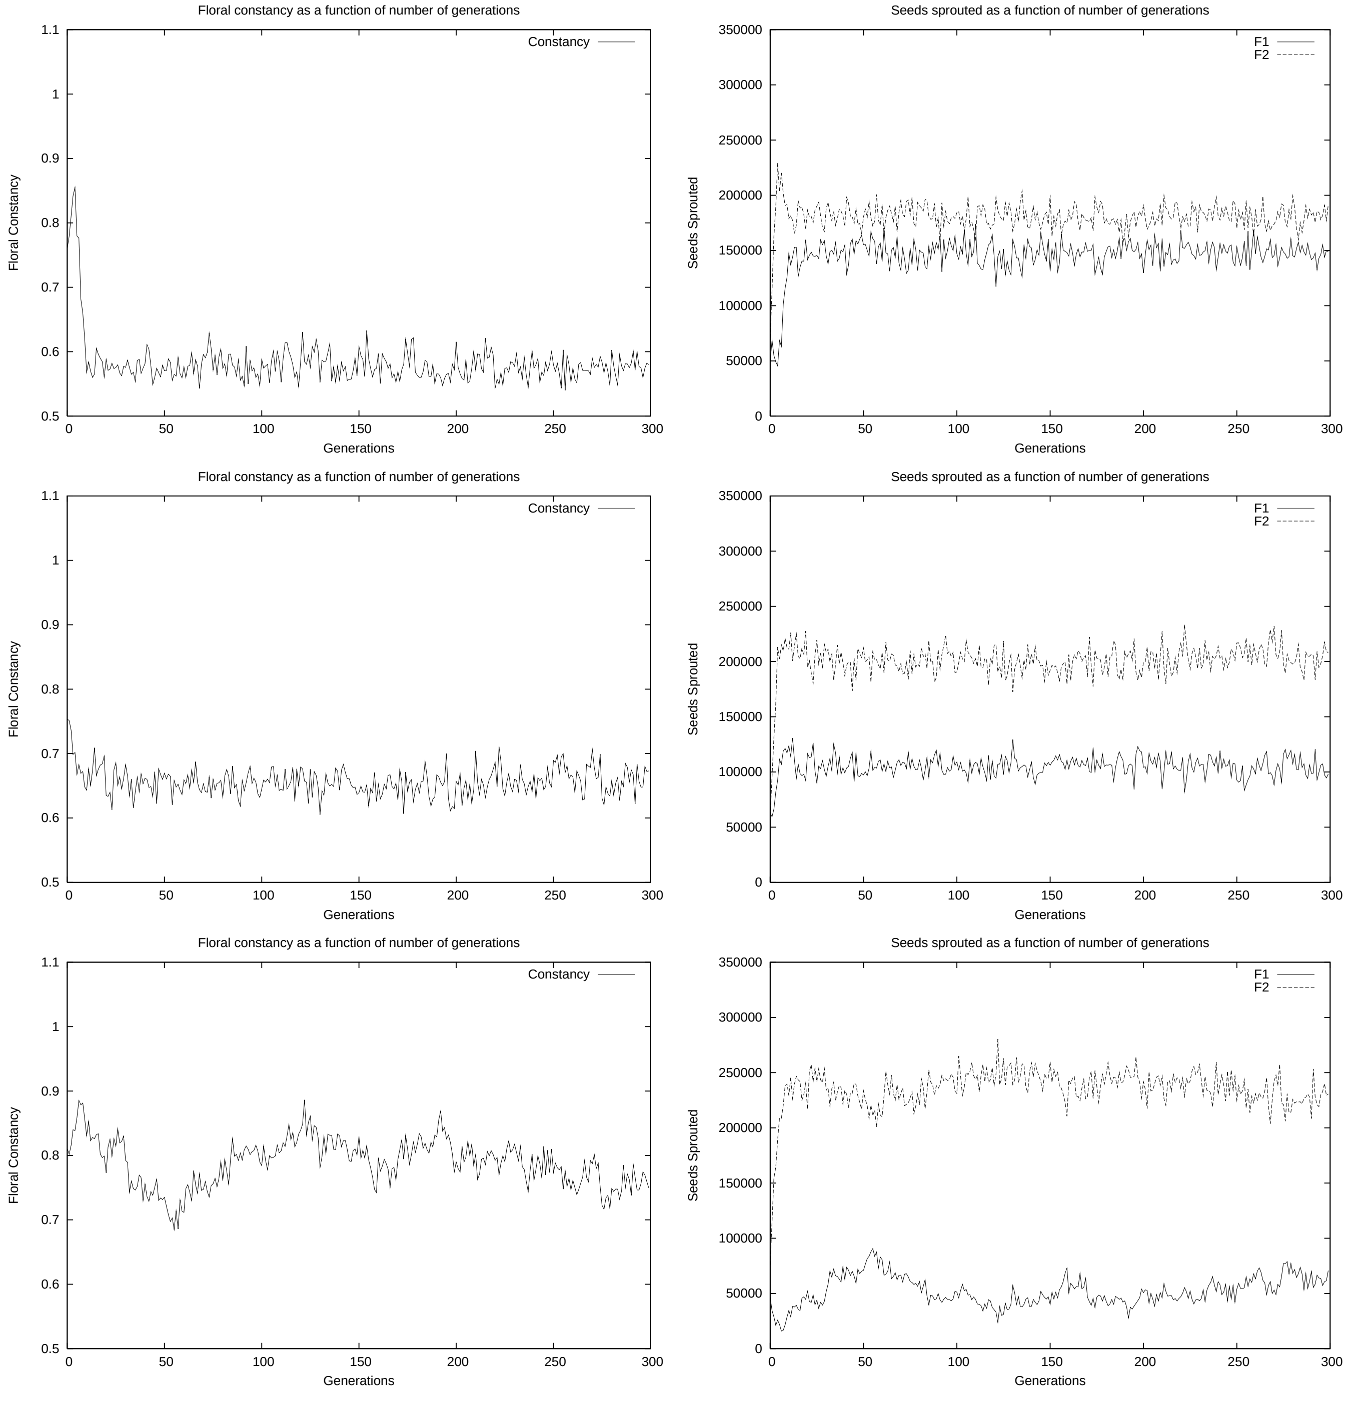

Figure 5: Floral constancy and seed sprouting levels for the three least specialized bees

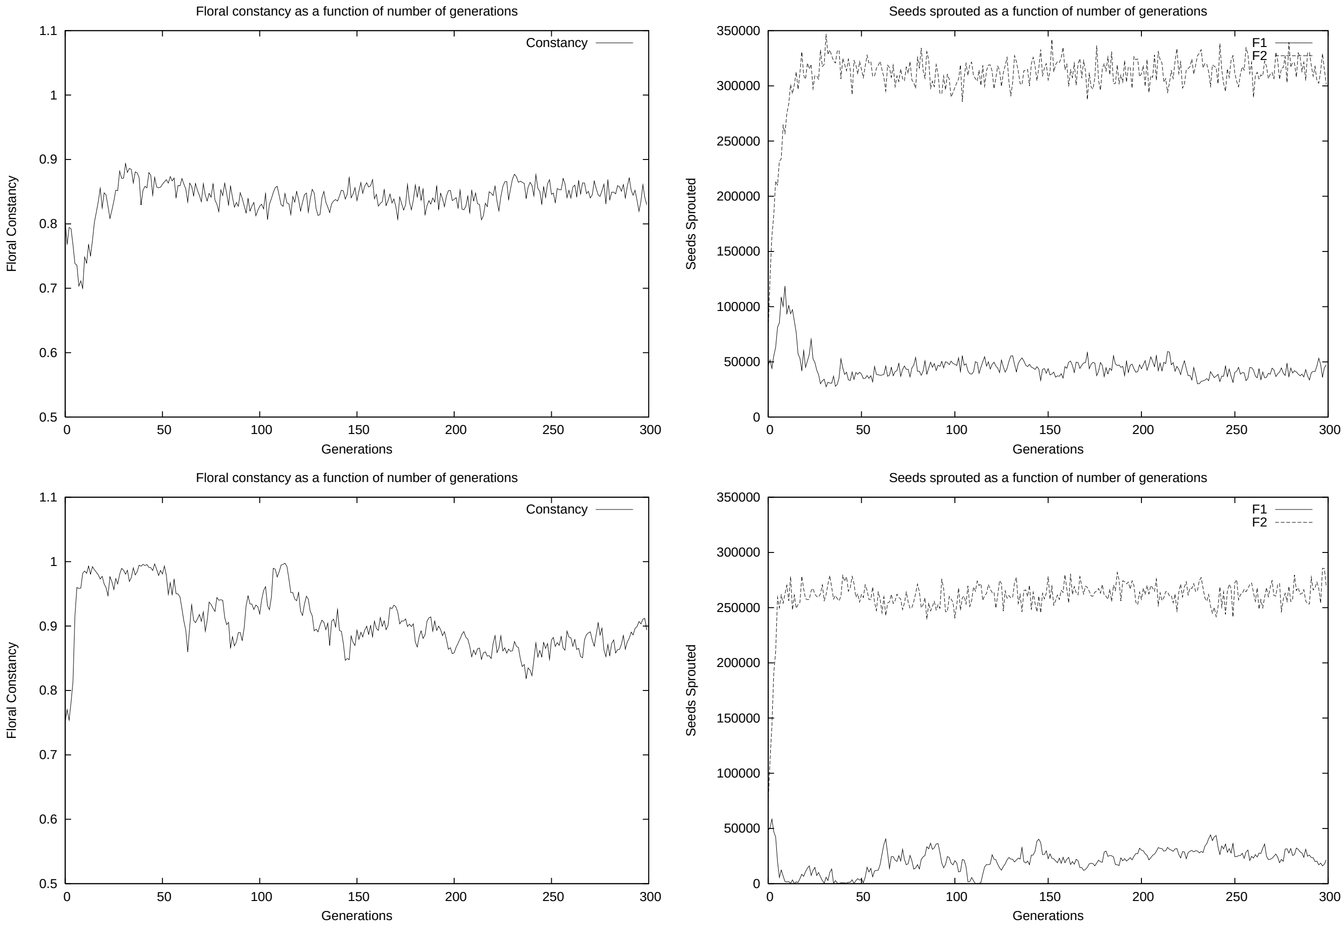

Figure 6: Floral constancy and seed sprouting levels for the two most specialized bees

# 4 Discussion and Next Steps

The results of this study indicate that the virtual bees in this study require substantial pressure to specialize. It is observed that the bees only specialized when they were presented with one flower with next to no nectar and one with a significant amount and were allowed to fly only short distances. It is worth noting that the selection pressure on the bees to specialize is indirect. Bees that specialize in a single flower type ensure a richer crop of that type of flower next year because they pollinate that flower type more efficiently. Extending the fitness evaluation of the virtual bees to six years (plant growing seasons) was intended to give this indirect fitness pressure some scope to act. The observation that a large differential in nectar levels between the two flower types is required to cause the virtual bees to specialize suggests that this indirect selection pressure is acting weakly.

The choice of three pollen particles per flower visit should be re-examined in this light. A bee that visits both types of flowers is likely to pollinate a fairly large number of flowers of both types when dropping three particles per visit. Having both sorts of flowers persist is thus not too difficult. Reducing the number of pollen particles per visit may increase the scope of the weak selection pressure resulting from the health of overall flower populations on the bees.

An initial increase in fitness level for all populations before leveling off was observed in the first data set. This length was close to the same length as the specialization curve. Figure 7 shows the fitness over evolution of a population of specializing and non-specializing bees. Compare this with Figure 3. These plots seem to indicate that learning how to gather resources and deciding to specialize or not are linked. This linkage may be tied to the emergence of a single dominant genotype within the population of virtual bees.

Another potential for further enhancement of the model is in adding costs to flight length and eating of nectar. The current model does not take into account the energy that must be expended during the course of gathering nectar. Factoring in these costs could result in populations needing to specialize more frequently in order to survive.

Further study is planned into the behavior and preferences of individual automata. Analyzing individual bees for preferences may explain some of the populations which came close to specializing but returned to a lower stable state. Heinrich's study of bee response to levels of nectar in white clover could be reproduced in the context of the virtual bees and would provide a good comparison of the automata with documented bee behavior. Also possible is a study involving both nectar quantity and quality to simulate situations where abundant watery nectar is less desirable than smaller amounts of better quality nectar.

Study of behavioral features such as traplining will require upgrading the current movement algorithm for the virtual bees. At present it is a random walk with only the decision to drink or not at randomly encountered flowers being evolved. This approach was chosen to ensure simplicity in the initial study.

The plant model used in this study is minimal. A flower could be harvested for nectar exactly once, whether it was successfully pollenated or not. Allowing for continued nectar production until successful pollenation would complicate the pressures exerted on the bees. Virtual bees would be faced with a situation where maximizing available food this year by not pollenating any flowers results in minimal nectar availability next year when none of the flowers reproduced.

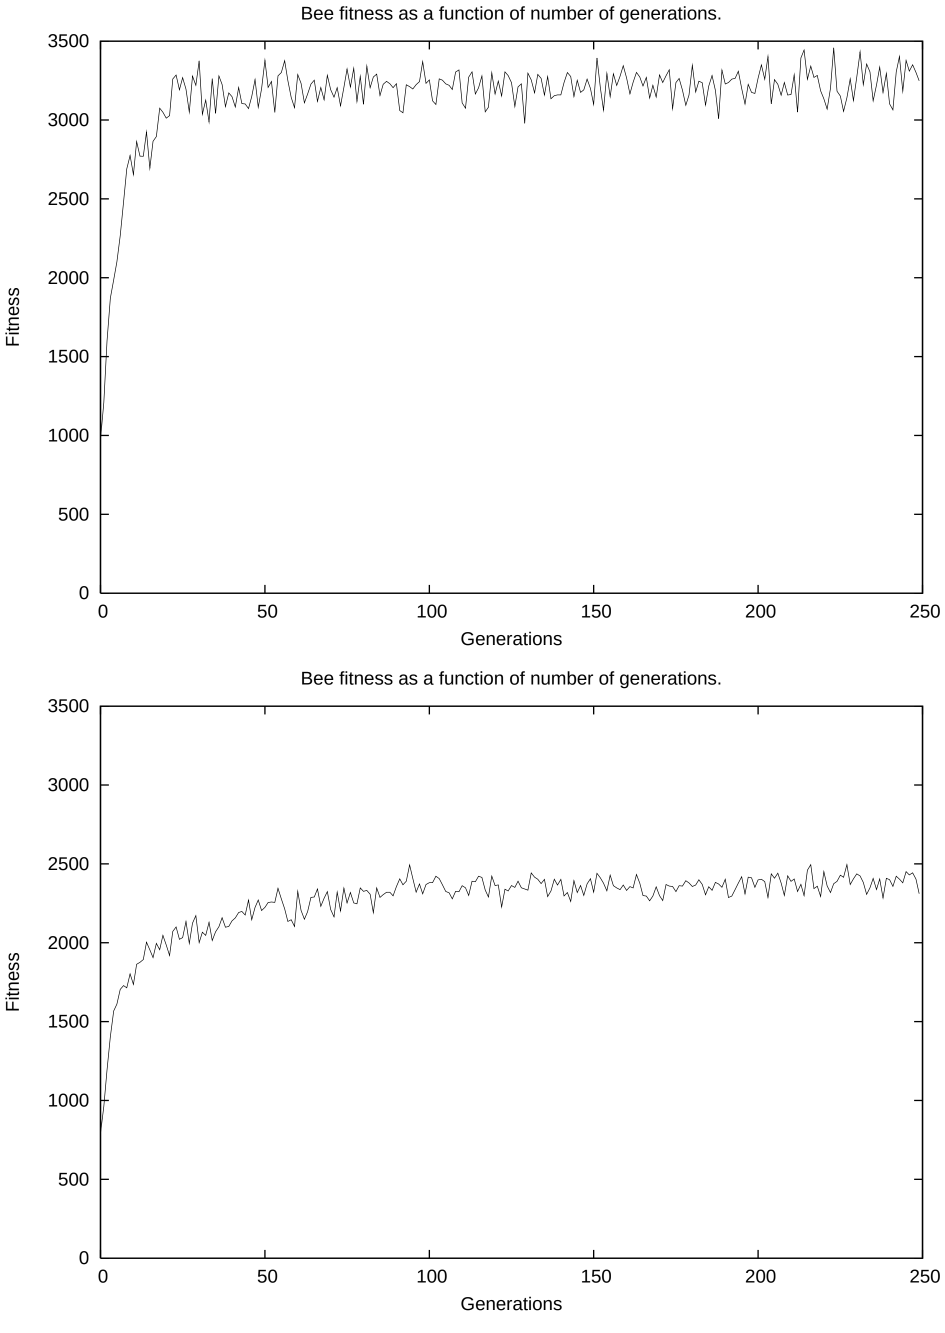

Figure 7: Fitness over the course of evolution for a population that exhibited floral specialization (top) and one that did not (bottom).

# 5 Acknowledgments

The authors would like to thank the Iowa State University Department of Mathematics for its support of the second author during the period this research was conducted. They would also like to thank Ralph Lewis of Swisher, Iowa for allowing observation of his bee keeping operation.

# References

[1] D. Ashlock. 2004. Evolutionary Computation for Modeling and Optimization.

[2] F. G. Barth. 1985. *Insects and flowers.* Allen and Unwin, London.

[3] C. Darwin. 1876. Cross and self-fertilization. in the vegetable kingdom. J. Murray, London.

[4] B. Heinrich. 1976. Foraging specializations of individual bumblebees. Ecological Monograph. 46:105-128.

[5] B. Heinrich. 1979. *Bumblebee economics.* Harvard University Press, Cambridge, MA.

[6] H. C. Lewis. 1986. *Memory Constraints and flower choice in* Pieris rapae. Science 232:863-865.

[7] C. D. Michener. 1974. *The social behavior of the bees: a comparative study.* Belknap Press of Harvard University Press, Cambridge, MA.

[8] G. H Pyke. 1978. Optimal foraging: Movement patterns of bumblebees between inflorescences. Theoretical Population Biology 13:72-98.

[9] D. W. Roubik. 1989. Ecology and natural history of tropical bees. Cambridge University Press, Cambridge, MA.

[10] P. Schmid-Hempel. 1984. *The importance of handling time for the flight directionality in bees.* Behavioral Ecology and Sociobiology 15:303-309.

[11] N. M. Waser. 1983. The adaptive nature of floral traits: ideas and evidence. Pp 242-286 in L. K. Real, ed. Pollination biology. Academic Press, Orlando, FL.

[12] M. Zimmerman. 1979. *Optimal foraging: A case for random movement.* Oceologia 43:261-267.