

# “Bees?” A Large Scale, Cooperative Simulation Weighing Altruism and Selfishness

ALEXANDER SIMMS, RAVENNA THIELSTROM

Swarthmore College

## Abstract

*In nature, communities that revolve around altruistic cooperation exist: for example, a bee colony, where individual bees retrieve nectar for the communal hive rather than scouting selfishly for their own survival. Although cooperation can certainly be implemented into an artificial life simulation, from an evolutionary standpoint a more interesting question is raised: can the interactions of a group and the parameters of their environment produce emergent cooperation among independent agents? This paper aims to investigate the possibility of emergent altruism through evolutionary game theory. To this effect, NEAT was used to evolve neural-net topologies across many generations in a bee colony simulation, testing what possible circumstances might lead to altruistic decisions or selfish decisions in the hive. A number of experiments were conducted, incorporating such concepts as a measure of group fitness, recurrent networks, and inter-agent communication. Initial speculation produced a hypothesis that the amount of altruism and selfishness demonstrated in NEAT-trained agents will be most affected by an individual fitness relying on group fitness. The results of the experiments demonstrate that while this is partly true and group fitness is key to altering artificial behavior, selfishness generally continues to prevail, with cooperation appearing rarely, or only when coerced. Through group-affected fitness, agents can learn to use altruism strategically to mitigate the negative effects of selfishness.*

## I. INTRODUCTION

Cooperative behavior has frequently been discussed within the contexts of artificial intelligence. The prisoner’s dilemma often serves as a basis for these discussions due to its relevance to issues and challenges of social order. Michael W. Macy’s paper on “Social Order in Artificial Worlds” describes the prisoner’s dilemma as the “paradox of social order”. [3] The prisoner’s dilemma is a hypothetical situation that tests human nature. The basic dilemma is this: Two criminals are imprisoned, awaiting trial. The chief prosecutor has enough evidence to convict both of them of a lesser charge, but not quite enough evidence to convict them of their full crimes. The prosecutor then hatches a plan: both prisoners are asked to betray the other by testifying against him. If one prisoner should choose to defect and betray the other, he will be

set free and the betrayed prisoner will receive 3 years in prison. If, on the other hand, they both simultaneously betray each other, both will receive a 2 year sentence. Cooperative silence from both parties gets them both the lesser sentence of 1 year. Although it may be argued that altruism is the best option for both parties, according to most game-theoretic explanations of this scenario, one should always choose to betray in order to minimize personal loss, and indeed one would expect this kind of decision from an artificial agent. Macy specifically names four possible outcomes in the prisoner’s dilemma, ordered by the amount of payoff: temptation, defecting while the other prisoner remains silent; reward, both prisoners staying silent; punishment, both prisoners defect; and “sucker,” staying silent while the other prisoner betrays you. The reason why this game can be called a paradox is that the choice leading

to the highest payoff (selfishness) also has the potential for the greatest failure. The experiments outlined in this paper take place not on an individual-to-individual level but rather in a large communal colony: they don't explicitly use the prisoner's dilemma. However, each agent has a similar choice of selfishness or altruism with somewhat similar outcomes, only on a larger communal scale. It should be noted that the experiments outlined in this paper also differ from the prisoner's dilemma in another significant way: these experiments focus not on how selfishness in one agent causes suffering in another, but rather how selfishness negatively impacts the entire community overall, including the selfish individuals themselves.

As mentioned by Macy, the single-iteration prisoner's dilemma is best played with a selfish dominant strategy, a conclusion that is also borne out in the experiments of this paper. However, this strategy is not the best when the prisoner's dilemma is repeated, as expanded upon in David B. Fogel's "Evolving Behaviors in the Iterated Prisoner's Dilemma". [1] In the iterated prisoner's dilemma, both prisoners are put through this same procedure time and time again, and have memory of how the iterations went in the past. Fogel explicitly evolved best strategies over many different simulations of the scenario. Players of the recurrent prisoner's dilemma will learn over time which individuals are more likely to defect or cooperate, and adjust their strategies accordingly when facing each of those individuals. However, in a hive environment where many independent agents are making the decision simultaneously, this is not a viable option to learn. Instead, for some of the experiments, recurrent networks are used to provide each generation of agents with several rounds of simulations and information from each one in order to learn general trends of the hive.

Finally, Sarit Kraus' "Negotiation and cooperation in multi-agent environments" discusses the possibilities of achieving cooperation in a shared environment like the hive simulation set

up for these experiments. [2] One such strategy discussed suggests that in situations where resources must be shared, a communication system may benefit self-motivated agents. Along a similar line of thinking, Section V of this paper implements a system in which agents have the opportunity of increasing overall hive resources by communicating where more resources can be found.

## II. EXPERIMENTS

### I. The Bee Model

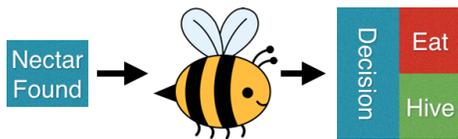
To conduct these experiments, a model was developed that takes inspiration from the nectar-collecting hive insect, the bee. Our model consists of a population of individual neural nets, evolving by the NeuroEvolution of Adapting Topologies algorithm, or NEAT. [4] NEAT, first developed by Kenneth O. Stanley, is a genetic algorithm that has a flexible genetic encoding that allows for the topology of the network to grow and change over time, based on a user-defined measure of "fitness". All of the genes within the genetic encoding are tagged with historical markings, which allows for efficient implementation of "crossing-over," the process by which genetic information from an individual's parents is mixed. Finally, NEAT allows for speciation: individuals within the population are grouped together based on the similarity of their genomes, and only those individuals with similar genomes are allowed to reproduce with one another. This allows for each species to find its own niche, and prevents any one species from outcompeting all of the others. It is because of these features that NEAT was selected for this task.

Each individual NEAT agent within this model is called a "bee," and the population of all of the NEAT agents is called the "hive". Each "day" within this model, all of the bees leave the hive to go and find nectar. The nectar is represented in this model as a random floating-point number between 0 and 1. When it receives this nectar, it has the choice to either be selfish or altruistic. If the bee decides to be

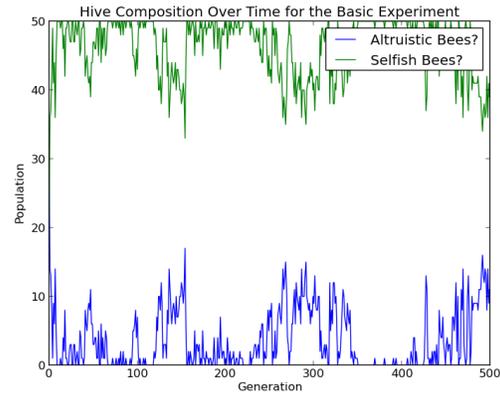
selfish, it eats the nectar then and there. If it decides to be altruistic, it brings the nectar back to the hive, where all of the nectar brought back by altruistic bees is pooled and shared equally among all of the altruistic bees. The “fitness” in this system is simply how much nectar a bee was able to eat on a given day.

## II. Basic Experiment

An experiment was carried out to determine the workability of the model. In this experiment, each bee lived for only one day, and afterwards died. Their fitnesses determined the makeup of the next generation of bees: the most fit bees were crossed to produce hopefully more fit offspring. The NEAT networks were evaluated using serial activation on a sigmoid function, as they were in all of the experiments detailed in this paper. The initial topology had no hidden nodes. Its only input was the nectar it found, again, a float between 0 and 1, and output represented the choice it made. (See Figure 1.)



**Figure 1:** All that is input into the system is the amount of nectar that the bee has found on that day, and the choice that the bee has made is determined by whether its one output is closer to 0 or to 1.



**Figure 2:** Under the conditions of the basic experiment, altruism lost handily to selfishness.

In all of the trials run, the majority of the bees wound up selfish. A representative trial is shown in Figure 2. There are a number of reasons for this, the most pressing being the fact that the bees had no direct incentive to act altruistically. While more stable fitnesses could theoretically be obtained through altruism, the system only works if all of the bees are acting altruistically: enough bees will bring back above-average amounts of nectar to offset the bees that bring back a below-average amount of nectar. However, if too few bees are acting altruistically, then this buffering effect is not present.

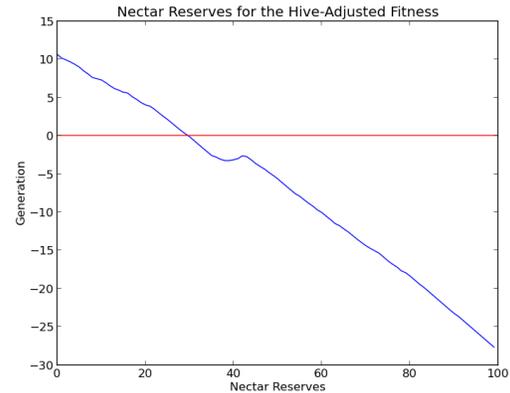
## III. Hive Fitness

How, then, can this selfishness be overcome? In the real world, actual hives collect the nectar that the bees bring in, and convert it into honey. This is the inspiration for a measure of “hive fitness,” indicating how much nectar the hive has on “reserve”. Each day, when the altruistic bees bring the nectar back, if nectar levels are below 10 units, a tenth of the nectar is taken from them and put into the hive, to support the “queen,” who eats 0.5 units of nectar per day. The amount of nectar in the reserve is turned into a modifier as described in Figure 3, which is then multiplied by the amount of nectar each bee finds to determine their fitness.

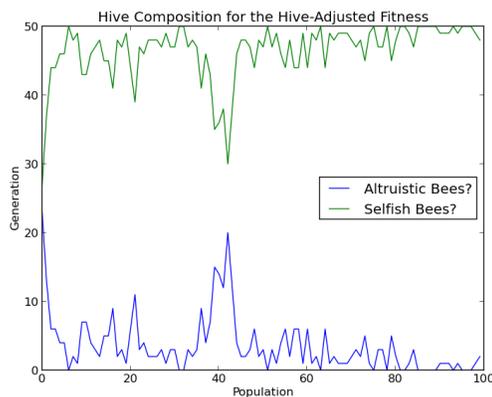
```

if hive_nectar > 10:
    modifier = 1
elif hive_nectar > 0:
    modifier = hive_nectar / 10
else:
    modifier = .00001
    
```

**Figure 3:** The fitness of each bee is affected by the amount of nectar reserves in the hive. The modifier defined by this section of code is multiplied by the amount of nectar the bee collected in order to determine that bee's fitness. The last case has a very small number and not just 0 due to restrictions in the NEAT algorithm.



**Figure 5:** Although the bees' fitnesses continued to shrink along with the hive fitness, they did not behave altruistically.



**Figure 4:** The addition of the hive fitness component to bee fitness was unsuccessful in coercing the bees to behave more altruistically.

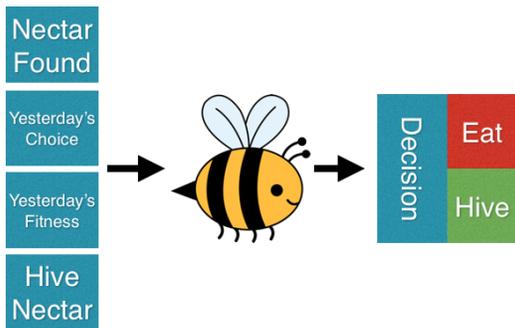
In this experiment, the bees still behaved overwhelmingly selfishly, for largely the same reasons as in the basic experiment. Although the hive fitness mechanic was added, the bees have no way of determining what the hive fitness currently is, and therefore have no way to adjust their actions to be altruistic or selfish, depending on the current state of the hive. Because of this, they continue to default to the same behavior of selfishness, because they are unaware of this new restriction on their behavior. Furthermore, because the same fitness modifier is applied to both the selfish bees and to the altruistic bees, the selfish bees have a similar advantage: 0.0007 is still better than 0.0001.

Subsequent experiments were run as a proof of concept, in which the fitness adjustment was only made to the selfish bees, and in which the selfish bees were all assigned a fitness value of 0. While these produced results in which the bees all wound up altruistic, further reporting on them is outside of the scope of this paper, which is attempting to get altruism without specifically encoding for it.

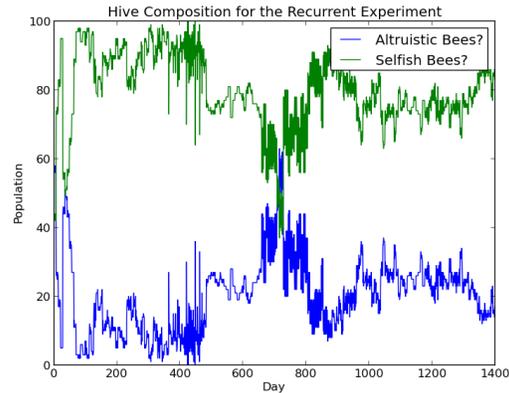
#### IV. Recurrent Networks

In this experiment, the bee model is considerably more fleshed-out. The bees now live for more than one day, and have considerably more inputs than they did before. Each individual day proceeds as it did in the last experiment: the bees go out to collect nectar, bring it back to the hive or eat it while they're out, and the nectar brought back to the hive is shared between the bees that brought back nectar, minus a 10% to support the hive itself, if nectar levels drop below 20. However, the bees now have a way of determining how much nectar is in the hive. This is a very important input, because their fitness rests on keeping this value within appropriate ranges. This input is normalized in the same manner that fitness is in Figure 3.

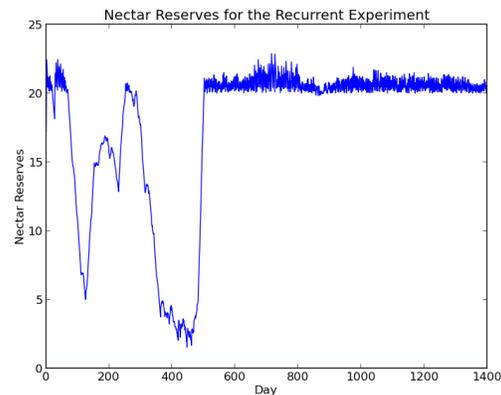
Recurrent networks can be much better at predicting the results that their behavior will have on their environment. Because of this, recurrent networks will be used in this experiment. To get the most out of this recurrence, the bees will also have inputs corresponding to the choice they made "yesterday," as well as how much fitness they got. (See Figure 6) On the first day of the bee's life, the recurrent inputs are set to 0.5 to avoid influences on the bee's behavior.



**Figure 6:** In addition to the nectar found, the networks now include the choice that the bee made yesterday, as well as the amount of fitness they received, and the level of the hive nectar, normalized between 0 and 1.



**Figure 7:** The wavering pattern shown in the later generations cropped up very consistently across experiments, and serves to maintain hive nectar at a particular level.

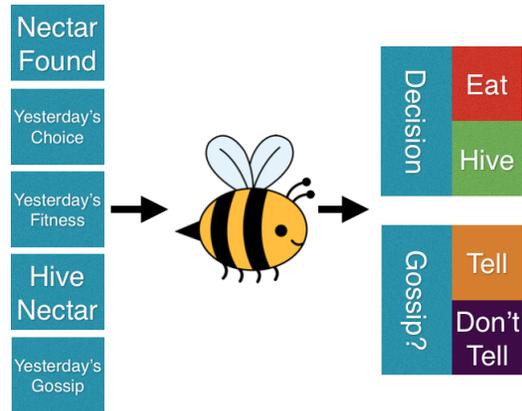


**Figure 8:** Under the recurrent network, the bees were able to maintain relatively high hive fitness.

Under the recurrent system, consistently higher hive fitnesses were maintained (see Figure 8), despite the fact that the majority of the bees were selfish (see Figure 7). This is because the bees are acting just altruistic enough to keep the hive alive, while acting selfishly for the rest of the time. This is facilitated largely by the input reminding the bees of hive nectar reserves. The nectar levels dip down twice in the experiment shown in Figure 8, however: these two drops represent

the networks learning the implications of that input on their fitness. After it is discovered, they maintain the hive fitness at a reasonable level.

Despite the generally positive hive fitness maintenance, overall it is clear that the bees favor selfishness. In all trials, altruism never rises substantially above the halfway point by the end of the specified amount of generations, except for a single instance which occurred during a series of runs with one starting hidden node. This will be discussed later with similar results from Section 2.5.

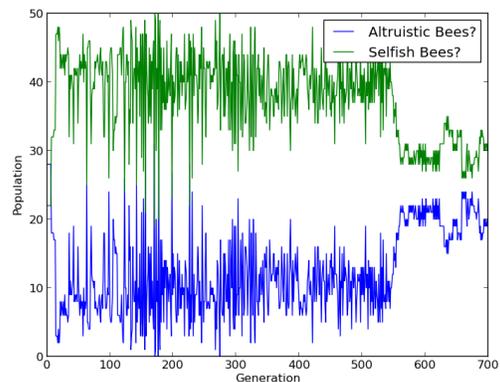


**Figure 9:** In addition to the first decision output, the second output represents a choice to share helpful information with other bees.

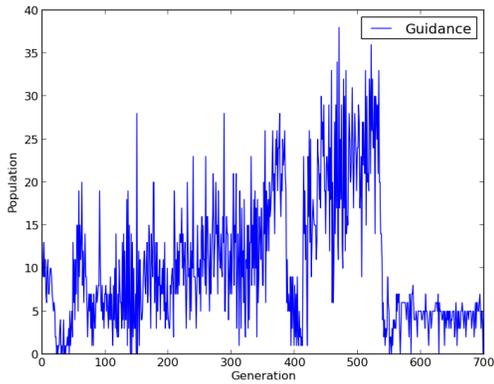
## V. Advice from Bees

One major factor that has been missing from this simulation of social order so far is, of course, direct interaction between the individuals. Each agent has control over their own decision, which may indirectly affect the destiny of the entire hive, but beyond this they have no opportunity to influence other members of their community, as does happen in real life. For example, another way that real biological bees cooperate within the hive is by sharing information about sites where they have foraged, essentially notifying other bees where to get the most nectar from.

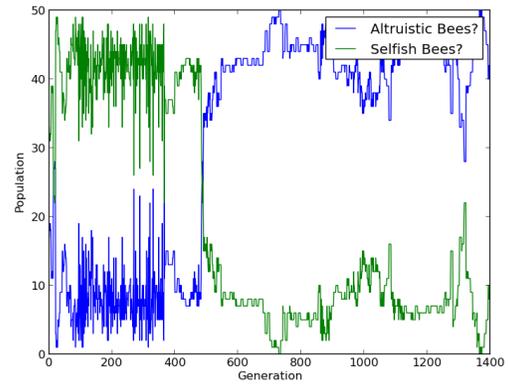
To implement this, a second output was added to the neural network: the decision whether or not to give another bee advice. (see Figure 9) This decision gets taken into account by the next bee in the population. If advice was given *and* the previous bee found enough nectar for the advice to be beneficial, the current bee will take that advice and receive the same amount of found nectar as the previous bee. Otherwise, found nectar will be randomized as usual. The general aim of this experiment was to observe how the bees would choose to exploit this opportunity and what trends in altruism or selfishness might result from prosperous advice.



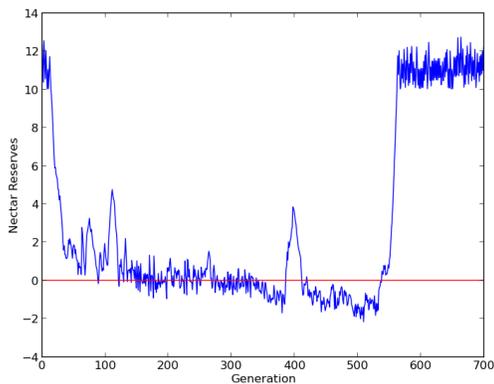
**Figure 10:** The advice system doesn't seem to generally improve upon previous experiments, but semi-positive instances certainly did occur. The familiar wavering was often present.



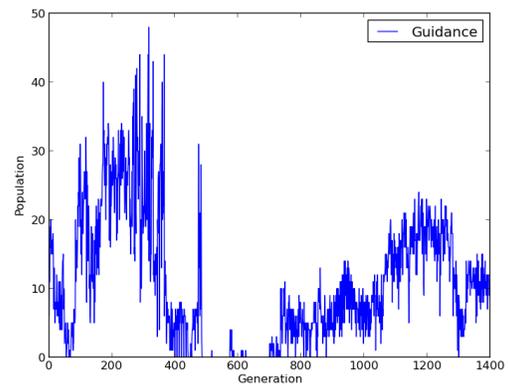
**Figure 11:** *The correlation between guidance and altruism was generally found to be inversely proportionate, although this wasn't always the case.*



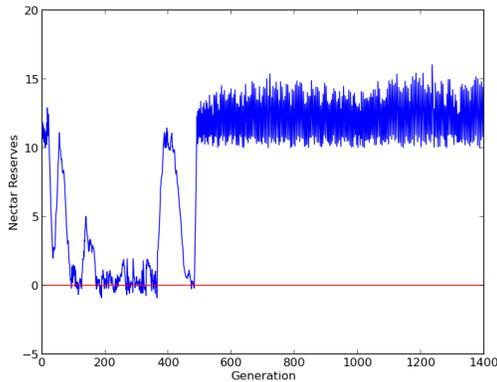
**Figure 13:** *A cooperative, altruistic society can be achieved very rarely.*



**Figure 12:** *Hive fitness also is seemingly inversely proportionate to guidance.*



**Figure 14:** *For altruism to win out, it seems guidance must be first abandoned, as can be seen around the 500 day mark.*



**Figure 15:** *As before, there seems to be a proportional relationship between guidance and selfishness.*

Results from this experiment were inconclusive. Increases in guidance within the bee population generally corresponded with increases in selfishness, as seen in the Figures 9 - 15 above. However, hive fitness was affected in two different ways: some runs showed increases in hive fitness as guidance rose, and some runs showed decreases (although the latter was generally more prevalent). The rationale behind this is that as guidance rises, the overall amount of found nectar rises, so hive fitness is more likely to rise if there are an acceptable amount of altruistic bees. On the other hand, bees that are consistently receiving an above average amount of nectar through guidance are also more likely to behave selfishly instead of altruistically, which obviously will decrease hive nectar. Overall, it seems that a lack of guidance actually corresponds to higher amounts of altruism, although to understand why this might be the case further research would be necessary. The system of one-way communication implemented here, then, seems to achieve the opposite of the desired affect.

Yet, for about 5% of the trials, evolutions were found that ended up with an altruistic majority. Experiments under the exact same conditions and parameters were also repeatedly run without the advice output influencing the found nectar, and exactly one

majority altruistic instance was found in a series of over twenty trials. Guidance, despite no longer having any affect on any part of the experiment, still showed the same patterns of correlation with selfishness.

### III. DISCUSSION

As has been demonstrated, it is rather difficult to get independent agents to achieve altruism in this sort of environment. It has been theorized that actual bees have some sort of urge hard-coded into their genetics to be altruistic, and this is how hives are able to survive. [3]

One substantial difficulty for the development of altruism in this particular model comes from the way that the fitness function is implemented. All of the bees coming back are coming back with nectar in the range of 0 and 1, so, amortized, choosing to be altruistic is choosing a fitness of 0.5. Choosing to be selfish has the potential to obtain much higher fitnesses than that, and those individuals will be able to reproduce, at the cost of the altruistic individuals of the hive.

It is important to note that altruism is of course perfectly possible and easily implemented, as has been mentioned in Section III, along with other test runs in which each bee's fitness was simply the overall hive fitness. If the programmer explicitly rewards or punishes behaviors, or directly links motivation solely to the group rather than the individual, good results can be achieved, but at the cost of the factor of self-motivation and inherent self-interest that is present in real life communities.

One phenomenon of the experiment which occurred again and again was the wavering behavior of the agents' decisions, demonstrated in Figures 7, 10, and 13. The larger the wavering is, the more likely it is to correlate to a stabilized amount of hive nectar, usually at 0, but with smaller oscillations the wavering can signify a higher stabilized level depending on the conditions of the hive fitness modifications. Wavering is caused by

the population rapidly varying the amount of altruistic/selfish decisions across days and generations. Since this generally occurs around one of the boundaries for punishment, it is logical to conclude that the bees, while not behaving truly cooperatively, can still evolve somewhat of a system for cooperation where altruism is practiced only in necessary moments to avoid hive punishment, and then reverting to selfishness once the danger has temporarily passed. This is obviously not an ideal community, especially since the oscillating tends to happen at 0, but the balance of selfish cooperation they have managed to achieve does grimly mimic some aspects of our own society.

Finally, the unexpected, rare appearance of unmistakable altruism during the advice system experiments is a result that cannot really be suitably explained with the knowledge available. The few instances where altruism appeared, as mentioned earlier, occurred only about 5% of the time in the advice system runs, and the only clear pattern between them seemed to be that guidance dropped down to zero or close to zero just as altruism began to win out over selfishness. Why this happens is unclear since there doesn't seem to be any benefit to not giving advice, other than avoiding the discussed trend of selfishness resulting from the higher amount of nectar found. The sudden increase in randomness could have had something to do with this behavior, as opposed to feeding the network growingly similar above-average values of nectar.

However, the fact that this pattern occurs even when guidance is made irrelevant seems

to imply that it is not the amount of nectar itself which prompts a higher chance of altruistic solution, but rather the presence of a second output which, since it is recursively passed onto the next day as an input, may serve as a kind of extra hidden node. This theory is supported by the single instance of altruism found in the recurrent network experiments, where 1 hidden node was added to the start of each experimental run. More experimentation would be needed to fully investigate the context of these rare instances.

Although results are not as clear or as optimistic as initially hoped for, the experiments conducted have shown that individual evolution of altruistic communities is indeed possible, if rare. Further research would be necessary to better understand how independent, self-motivated cooperation is formed in these instances.

#### REFERENCES

- [1] David B. Fogel. Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation*, 1(1):77–97, 1993.
- [2] Sarit Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94:79–97, 1997.
- [3] Michael W. Macy. Social order in artificial words. *Journal of Artificial Societies and Social Simulation*, 1(1), 1998.
- [4] Kenneth Stanley. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21, 2004.