

Evolving Agent Swarms for Clustering and Sorting

Vegard Hartmann

Complex Adaptive Organically-Inspired Systems Group (CAOS)

Dept of Computer Science

The Norwegian University of Science and Technology (NTNU)

Trondheim, Norway

vegarhar@idi.ntnu.no

ABSTRACT

Social insect societies consist of individuals with simple behavior, yet at the collective level the societies are capable of solving complex tasks. These tasks span a wide range, including clustering, patch sorting, and annular sorting. It has proved particularly difficult to recreate the social insects ability to perform annular sorting, even for engineered solutions. In this paper we present the results we have obtained from evolving swarms of agents that are capable of performing clustering, patch sorting, and annular sorting. Most noteworthy is our solution to the 3-type annular sorting problem which has previously not been solved successfully.

Categories and Subject Descriptors: I.2.m [Miscellaneous]

General Terms: Experimentation

Keywords: Annular sorting, ant, clustering, genetic algorithm, neural network, sorting, swarm intelligence

1. INTRODUCTION

In social insects complex collective behavior emerges from the interactions of individuals that each performs simple behaviors. The collective behaviors span a wide range of tasks from foraging and nest construction to thermoregulation and brood sorting [3, 4]. Cemetery formation and brood sorting are two prominent examples of insects collective behavior. Cemetery formation has been observed in different ant species such as *Pheidole pallidula*, *Lasius niger* and *Messor sancta* [3, 5]. When ants die they are carried out of the nest by workers and deposited in piles outside the nest. The underlying mechanism for this clustering behavior appears to be an attraction between dead bodies mediated by the workers carrying the bodies. Small clusters of bodies attract the workers to deposit bodies, and through this positive feedback mechanism the clusters increase in size. Brood sorting has been observed in the ant *Leptothorax unifasciatus* [5, 6]. Deneubourg et al. [5] describes this behavior: "... the eggs

are arranged in a pile next to a pile of larvae and a further pile of cocoons, or else the three categories are placed in entirely different parts of the nest". The sorting of brood in the *Leptothorax* ant when moving to a new nest site has later been studied by Franks and Sendova-Franks [6]. Their studies show that the brood items are arranged in concentric annuli: "The standard pattern is for eggs and micro-larvae to be in the middle, with the larger larvae further from the center in order of increasing size. Pre-pupae and pupae are distributed in positions between the outer ring of the largest larvae and those of the next largest size" [6].

The collective behaviors of social insects have inspired computer scientists to perform both computer simulations and robotic experiments that attempt to replicate this behavior. There are two main reasons for this. The first is that the underlying mechanisms responsible for the behaviors are yet unknown. By replicating the behavior of the insects the underlying mechanisms may be found, and a better understanding of nature may follow. The other reason is that the behavior of social insects has many attractive features such as robustness and reliability through redundancy. Application of social insect behavior to computer science may thus lead to better techniques. Computer models based on the clustering and sorting of insects can lead to better performance in areas such as search, data mining, and experimental data analysis. An annular sort applied to the results of a search may for example aid in the presentation of the results by organizing them in concentric rings with the most relevant results closer to the center of the structure.

2. RELATED WORK

The clustering of dead bodies by ants inspired Deneubourg et al. [5] in their seminal article on clustering by a group of homogeneous agents. Their model contains ants that move about randomly, and is based on the principle that when ants discover an object, the probability of picking up the object is higher if the object is isolated. In a similar manner, the probability of dropping an object is higher if the ant is in an area containing many other objects. The basis of the model is thus two probabilities, one for picking up objects, and the other for dropping objects. The work performed by Deneubourg et al. [5] has inspired most of the later research in computer and robotic implementations of swarms of agents that perform clustering or sorting. A robotic implementation of clustering has been done by Beckers, Holland and Deneubourg [2]. In their experiment they use simple robots that do not have memory, can only decide if they are moving more than a specified number of objects,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'05, June 25–29, 2005, Washington, DC, USA.
Copyright 2005 ACM 1-59593-010-8/05/0006...\$5.00.

and that can only sense the local density of objects as being above or below a fixed threshold. The robots only have three different behaviors, with only one being active at any one time. Using these minimally equipped robots Beckers et al. [2] are capable of displaying clustering behavior by the robots.

Deneubourg et al. [5] also presents an extension of their clustering mechanism that sorts two different types of objects. The basis of this mechanism is also differences in probabilities between the picking up and dropping of objects. Objects that are isolated or located in an area with a high density of the other object type have a higher probability of being picked up. Likewise, the probability of dropping an object is higher if the ant is in an area with a high density of the same type of object. Lumer and Faieta [9] have generalized Deneubourg et al’s [5] model to apply it to exploratory data analysis. In contrary to Deneubourg et al’s [5] model, Lumer and Faieta [9] use a continuous similarity function to judge the similarity between different objects. They have also left the idea of having a homogenous population of agents. Sorting of objects has also been performed in robotic experiments [8, 10, 11]. In their work, Melhuish et al. [10] showed that it is possible to sort two types of objects using minimally equipped robots with a small rule set. The robots have a U-shaped gripper that allows the robot to push a single Frisbee. If the robot is pushing a Frisbee and collides with another Frisbee, the pushed Frisbee is deposited and the robot reverses. A simple two object segregation is achieved by varying the distances the two different Frisbee types are pulled back before being released. This enables the robots to perform a simple two-object sort. The pullback algorithm has later been extended by Melhuish et al. [11] to allow for sorting of more than two types of objects.

The most complex sorting behavior performed by social insects is annular sorting, ie. sorting into a target-like structure of concentric multi-type rings. Wilson et al. [14] have attempted to simulate this annular sorting behavior of the *Leptothorax* ant using minimalist robots. In their work they test if three different mechanisms are capable of forming an annular structure from different Frisbees. The first mechanism is investigated in a computer simulation and “. . .relies on object size and segregation combined with a simple clustering algorithm” [14]. The robots push objects, and when they hit an object while pushing another object, they will reverse and thereby leave the previously pushed object next to the other object. In their next experiment they achieve segregation of equally sized objects by varying the pullback distances before dropping the different objects. When the robots hit an object, they will pull the currently pushed object backwards for a specified distance depending on the type of the object, before dropping it. In their last experiment they introduce the concept of a ‘combined leaky integrator’ to adapt the pullback distances of the different types of objects.

The limited amount of work that has been done attempting to recreate the annular sorting performed by ants is an indication of how difficult the problem is. According to Holland and Melhuish [8], previous to their work there did not appear to be any accounts of sorting analogous to brood sorting being carried out by physical robots. Work on computer simulations of brood sorting have also been very limited to this date. Another indication of the difficulty of annular sorting is that the underlying mechanism is still

unknown [13]. Despite this, there have been varying suggestions as to what may cause the annular structure in brood sorting. One is the muesli effect [1] which is self-sorting according to size where small particles percolate to the bottom of the packet. In brood sorting the smaller particles could move between the larger brood and thus form a cluster in the center of the structure. Another suggestion is mechanisms based on differential adhesion, and Wilson et al. [14] have showed that this can be done with two types of objects. However, brood sorting in *Leptothorax* ants include five types of objects and Sendova-Franks [13] have therefore found it unlikely for models based on differential adhesion to be satisfactory. Franks and Sendova-Franks [6] suggest that deliberate spacing using pheromones or metabolic waste products may be the underlying mechanism, or the conditional probabilities used by Deneubourg et al. [5] combined with packing rules. Sendova-Franks et al. [13] have recently discovered that the brood sorting by *Leptothorax* has two distinct phases of the direction of brood movement. In the first phase the brood items are moved in a direction away from the nest entrance, and this phase constitutes a clustering phase. In the second phase the ants move the brood in a random direction and in this phase larger brood items diffuse outwards more quickly than lighter brood items.

3. METHOD

This study looks into the behavior of swarms of homogeneous agents called *ants* that are capable of clustering and sorting objects. These swarms of ants reside in an environment consisting of objects that may be picked up and carried. The environment is a toroidal grid with periodic boundary conditions. Ants that step out of the environment on one side will enter the environment on the opposite side. A grid in the environment may either contain one ant, one object, or both one ant and one object. An environment therefore consists of two parts. The first is a collection of objects that in the beginning are randomly dispersed throughout the environment, and as time goes by are moved by ants so that they form certain patterns. The second component of the environment is a swarm of ants that move around and pick up and deposit the objects. The environment has discrete time steps, and for each time unit all the ants are run in a random order. An important characteristic of the environment is the relationship between the size of the environment, the number of objects, and the number of ants in the environment. As suggested in previous work [7, 12] the number of objects can be used to decide the size of the world and the number of ants to be used. Ramos and Merelo [12] suggest that the number of ants, n_a , divided by the number of objects, n_o , should equal 0.1. When it comes to the size of the world we follow Handl et al’s [7] approach. They suggest that the length of each side in a square world should be $\sqrt{10 \times n_o}$.

The ants in the swarm are all identical, and they may perform three types of basic behavior: 1) pick up an object, 2) drop an object, and 3) move. An ant may carry only one object at a time. If an ant is not currently carrying an object it may attempt to pick up an object. The ant can only pick up an object that is located in the same grid in the environment as the ant itself. The second type of behavior that the ant may perform is dropping objects. If the ant is carrying an object, then it may drop this object at its current location if this does not already contain an




Figure 1: Mapping from an ant’s neighborhood to its neural controller.

object. The third and final behavior is movement. The ant has eight neighboring cells, and it may move to any one of these. However, an ant cannot move to a cell already occupied by another ant although it can move over objects.

Whether an ant should attempt to drop or pick up an item, as well as the type of movement it should perform is decided by a neural network controller. This network is a simple feed forward network with one hidden layer. The ant perceives the objects that are located in its eight neighboring cells, as well as the object it is currently carrying. These nine perceptions are the input to the network’s nine groups of input nodes. The mapping from the ant’s neighborhood to the input groups of the network is shown in Figure 1, where the input to group A is the object being carried by the ant. The input groups of the network contain one node for each object type in the environment, and only one node in each group may fire at any one time. To simulate the non-determinism of real-life behavior, noise is added to the perceptual input. In one percent of the perceptions the actual perception is randomly replaced by one of the other possible perceptions. The network has six outputs that control the behavior of the individual ants. The first two outputs (P and D) decide whether the ant should attempt to pick up or drop an object. The next two outputs (F and B) specify if the ant should move forwards or backwards, and they may cancel each other out. Each ant has a direction with eight possible values, and the last two outputs (L and R) specify whether this direction should be shifted left or right. The ants do however have one important simplification in that they are unable to perceive each other. This means that they may end up in deadlocks by wanting to move to another ant’s current location. To avoid this, an ant that remains in the same position for ten consecutive steps is moved to a random unoccupied position in the environment.

As described above an environment contains objects and a swarm of ants. An environment is equal to an individual in the evolutionary context. In the evolutionary run there is a generation of environments that each are run for a specified number of steps. When all the environments in the generation are completed, they are evaluated, and the next generation of environments is created. The phenotype of the environment is the positions of the objects, excluding objects that are currently carried by ants, after this has gone through the specified number of steps. The positions of the objects are used by the fitness function to evaluate the solution. The genotype of the environment is an array of real-valued numbers that define the weights in the neural network controller of the ants in the environment.

Many of the parameters are identical for all the experiments. All experiments have a population size of 50 that each run for 4000 steps, and the evolutionary run is 6000 generations long. The individuals each contain 60 or 50 objects, and this decides the number of ants and the size of the world. If there are 60 objects there are six ants in a 25×25 grid world, and if there are 50 objects there are 5 ants in a 23×23 grid world. Single point crossover is employed, and it can only happen between weights that belong to different nodes, that is, in crossover all weights applied to the inputs of the same node are transferred to the next individual. The crossover and mutation rates are constant for all the experiments with the crossover rate being 0.7 and the mutation rate 0.05. Tournament selection in combination with elitism is used as selection mechanism, and each trial contains 5 individuals. The chance of choosing the best individual is 0.75, and if this is not chosen, a random individual from the trial is chosen for reproduction. The 5 best individuals from each generation are copied into the next generation only subject to mutation.

4. EXPERIMENTS AND ANALYSIS

We have performed four experiments with different levels of complexity. The simplest one is what Melhuish et al. [10] term clustering. This is the task of grouping objects of the same type in a continuous area. The next task is to perform patch sorting [10]. This involves two or more classes of objects that are individually clustered, while the clusters of each class are separated from each other. The last two experiments are considered to be of a higher complexity. These tasks are annular sorting of two and three classes of objects. Melhuish et al. [10] defines annular sorting as “forming a cluster of one class of objects, and surrounding it with annular bands of the other classes, each band containing objects of only one type”. Thus the two tasks are to form a central cluster surrounded by a single band of another class of objects, and forming a central cluster surrounded by two bands of two other types of objects, each band containing only objects from the same class.

4.1 Clustering

The first experiment attempts to solve the problem of clustering identical objects. Here all the objects must be grouped together such that they occupy a continuous area. This is equivalent to cemetery formation in ants. To solve this problem we have created a local fitness function based on the idea that each object should be surrounded by as many other objects as possible. The fitness function is given in Equation 1. In the equation, n_o is the number of objects




Figure 2: A single cluster is formed after 5840 steps.

in the world and o_i is an object. The function $c(o_i)$ returns the number of objects that are located in object o_i 's eight cell neighborhood.

$$\sum_{i=1}^{n_o} c(o_i) \quad (1)$$

To control the individual ants a feed forward network with three hidden nodes is used. The number of hidden nodes was determined experimentally, and it is one hidden node more than there are sensory stimuli. That is, the ants may sense an object or nothing. The individuals in the evolution contain 50 objects and 5 ants.

After the evolution has run for 6000 generations, the average fitness of the generation is 293.32 with a standard deviation of 52.44. The best individual has a fitness of 324 which is the maximum achievable fitness. As can be seen in Figure 2, this individual creates a single cluster with a perfect score in 5840 steps.

The clustering performed by the final solution appears to be based on two specific features of the ants picking up and deposition behaviors. First, the ants mostly pick up objects that are located in sparsely populated environments, and rarely remove objects that are part of larger, dense clusters. Second, the ants usually deposit their objects only in densely populated areas, although they may occasionally deposit an object next to pairs of objects. These characteristics of the ants' picking up and deposition behaviors enable the final solution to cluster a set of identical objects.

To evaluate the stability of the evolved solution, it was run for 10000 steps on 50 different random configurations containing 50 objects. Over these runs an average fitness of 319.44 with a standard deviation of 9.71 was achieved.

4.2 Three-Object Patch Sorting

The problem to be solved in this experiment is that of patch sorting [10]. In this experiment there are three different types of object. A successful solution must therefore create three different clusters that are separated, and that each contains only one type of object. To solve this problem we modified the fitness function used for simple clustering (Equation 1). Since there are now different types of objects, each object should not only be surrounded by as many objects of the same type as possible, but it should also have as few objects of a different type in its neighborhood as possible. The fitness function used is given in Equation 2. In the

equation $c(o_i)$ is a function that returns the number of objects in object o_i 's neighborhood, and $s(o_i, o_j)$ is a discrete similarity function that returns 1 if the objects are of the same type, and -1 if they are of different types.

$$\sum_{i=1}^{n_o} \sum_{j=1}^{c(o_i)} s(o_i, o_j) \quad (2)$$

The feedforward network used as controller for the ants has five hidden nodes. The reason for having four hidden nodes is the same as for the clustering problem. There is one more hidden node than there are sensory stimuli, which is three different object types plus nothing. The individuals in the evolution each contain 60 objects with 20 objects of each type, and 6 ants.

After the evolution has run for 6000 generations, the average fitness of the final generation is 270.88 with a standard deviation of 32.97. The best individual has a fitness of 314, which corresponds to 94% of the maximum fitness. As can be seen in Figure 3, the best individual is capable of forming three separate clusters after 7970 steps.

When the final solution is run on a three-object patch sort task, its ants quickly become loaded with mostly type 2 and 3 objects (empty circles and triangles in the figure). The reason for the ants not carrying type 1 objects (filled circles) is that these are deposited more easily in sparsely populated areas than the other two types of objects, and when picked up by the ants they are quickly deposited. Soon all the type 3 objects are part of groups, and one or two of these then quickly increase in size as the other groups of the same type are broken down. As this happens most of the type 2 objects also begin to form small groups. One of these groups then begins to grow, and at the same time the type 3 objects are all grouped into one cluster. If there were initially only one type 3 group that increased in size this will eventually contain all the objects of this type, but if it was two that initially began to grow, one of these will increase in size while the other is broken down. As all the type 3 objects become part of one large cluster and one or two of the type 2 groups begin to grow, most of the type 1 objects are also grouped into small groups. When this occurs the ants are not loaded as much of the time as they were earlier in the run of the individual. The run continues with the type 1 objects being collected in fewer, larger groups. At the same time the type 2 objects are all collected in a single cluster. Eventually all the type 1 objects also form a single cluster and at this point the objects have been sorted successfully.

When evaluating the stability of the solution over 50 different configurations of 60 objects, running for 10000 steps each, an average fitness of 279.68 with a standard deviation of 16.4 is achieved.

The evolutionary progression leading to the final solution evolves from solutions that cluster one type of object, to a solution that clusters all three types into separate clusters. At the beginning of the evolution only type 3 objects are clustered. The type 2 objects occasionally form small, fragile groups, while the type 1 objects remain randomly distributed in the world. The next evolved solution (generation 253) start by clustering the type 3 objects. When most of the type 3 objects become part of the cluster the ants begin to form groups of type 2 objects. These groups are however very fragile and are continuously formed and broken down. As with the first evolved solution, the last type of objects are




Figure 3: The formation of three separate clusters.

hardly moved at all and they remain randomly dispersed in the environment. After another 113 generations, the fittest individual is capable of clustering type 2 and 3 objects into one cluster of each type. However, the third type of objects is still randomly distributed in the environment as they are only carried one or two steps by the ants before being deposited. After a total of 3273 generations a solution that also carries the type 1 objects for significant periods of time is evolved. In this solution the ants first cluster the type 3 objects. As this cluster emerges the ants start forming groups of type 2 objects. All but one of these groups are after some time broken down, leading to the formation of a cluster of type 2 objects in addition to the cluster of type 3 objects. Once most of the type 2 and 3 objects become part of the two clusters the ants also begin clustering type 1 objects. These objects are however clustered into a loose cluster that is only barely connected. When looking at the evolving solution it is evident that the ants of the fittest individuals first learn to cluster type 3 objects. Then they learn to cluster type 2 objects, and finally a solution is evolved where the ants are capable of clustering all three types of objects.

4.3 Two-Object Annular Sorting

The task in this experiment is to create an annular structure containing two types of objects. The structure will thus consist of a central cluster of one object type surrounded by a band of the other object type. To achieve this task we have used a fitness function that is a modified version of the metric used by Wilson et al. [14] to judge the quality of annular structures. The fitness function consists of three of the four components included in the metric. The first component is the compactness metric. We have replaced Wilson et al’s [14] compactness metric with a function using the equation given in Equation 1. This metric takes the score obtained by Equation 1 and gives this as a fraction of the maximum achievable score with n_o objects in the world, $m(n_o)$.

$$C = 100 \times \frac{\sum_{i=1}^{n_o} c(o_i)}{m(n_o)} \quad (3)$$

The second component is the separation component. This component counts the number of objects that infringe on the ‘home zone’ of another object type. To do this the distance to the center of the structure is measured for each object,

and the upper and lower quartiles are computed for each object type. Then three different counts are performed. First, the number of central type objects that have a distance to the center greater than the lower quartile range of any other type is counted, N_c . Second, the number of outermost objects that have a distance to the center lower than the upper quartile of any other type is counted, N_o . Third, for intermediate type objects two types of objects are counted. These are objects that have a distance to the center greater than the lower quartile of any object type further from the center, N_i^g , and objects that have a distance to the center less than the upper quartile range of any object type closer to the center, N_i^l . These counts are then combined as shown in Equation 4.

$$Se = 100 \times \left(1 - \frac{N_c + N_o + \frac{N_i^g + N_i^l}{2}}{n_o} \right) \quad (4)$$

The last component that we use is the shape component. This is composed of two parts. The first is the fraction of central type objects that are located in the center of the annular structure, f_c . The second part calculates the average Euclidean distance from the center of the structure for each non-central type, \bar{d}_c , and sums up the deviations from this distance for all the objects of the same type, d_c . This summed deviation is then normalized. Equation 5 shows how the shape metric is computed, with m being the types of different objects.

$$Sh = \frac{100 \times f_c + \sum_{c=2}^m \left[100 \times \left(1 - \frac{d_c}{\bar{d}_c} \right) \right]}{m} \quad (5)$$

Influenced by the work of Sendova-Franks et al. [13] which found that the brood are first clustered and thereafter spread out to form the annular structure, we have given the compactness component twice the weight of the other two components when they are added together (Equation 6).

$$F = 2 \times C + Se + Sh \quad (6)$$

The feedforward network used to solve the task has four hidden nodes. Each individual contains 50 objects with 25 objects of each type, and 5 ants.

The final generation of the evolution has an average fitness of 362.48 with a standard deviation of 31.93. The best individual has a fitness of 392.54, which is 98.14% of the maximum achievable fitness. As can be seen in Figure 4, this

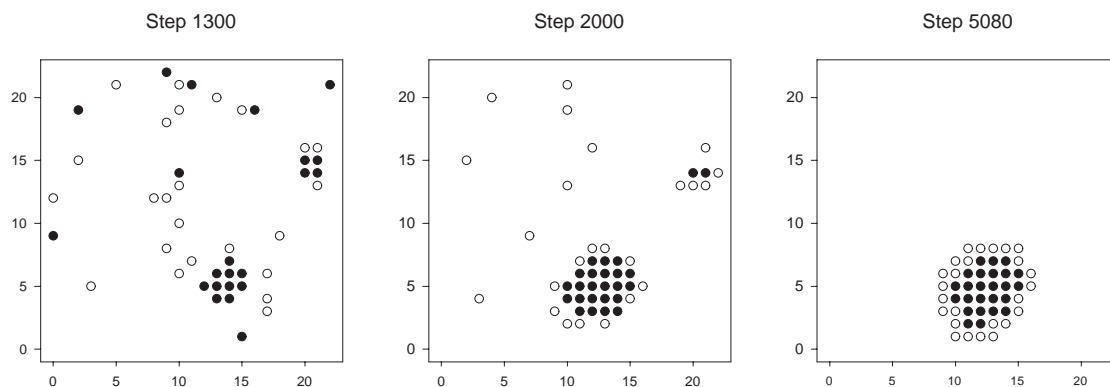


Figure 4: An annular structure with a fitness of 392.90 is created after 5080 steps.

individual creates a two-object annular cluster after 5080 steps. This structure has a fitness of 392.90, with a score of 100 for the compactness component, 100 for the separation component, and 92.90 for the shape component.

The cluster is created by the ants as follows. The ants quickly become loaded with an object of either type, but after a brief period the ants are only carrying central type objects. The reason for this is that these objects are only deposited in areas with a high density of central type objects, and there are no such areas in the world at the beginning of an annular sort. Ants that pick up central type objects thus become stuck with their carried object, while ants that pick up outermost type objects are able to deposit these again. After some time there is a fluctuation in the world that causes the formation of a small group of central type objects. This group quickly increase in size, and once a central type object become part of this group the ants rarely pick up the object. This causes the ants to carry the outermost type objects more of the time than previously in the run. These outermost type objects are for the most part put down along the edges of the central type cluster, but some are put down next to other outermost type objects or in isolated positions. However, the outermost type objects that are positioned along the edges of the cluster are rarely picked up by the ants. The ants mainly pick up the outermost type objects that are scattered around the world in small groups or in isolated positions and deposit these on the edge of the cluster. This cluster eventually contains all the objects and has an annular structure.

The solution that is evolved is quite stable, and when evaluating it on 50 different start configurations that each run for 10000 steps an average fitness of 360.78 with a standard deviation of 30.66 is achieved.

4.4 Three-Object Annular Sorting

The problem to be solved in this experiment is the same as for the two-object annular sorting experiment, only this time there are three types of objects. This means that there will now be an intermediate band of objects between the central cluster and the outermost band of objects. The fitness function used to solve the task is the same as that used for two-object annular sorting (Equation 6). The number of hidden nodes in the feedforward network and the number of ants and objects are the same as in the three-object patch sort experiment. This means that there are 5 hidden nodes

in the network, 6 ants, and a total of 60 objects with 20 objects of each type.

After 6000 generations an average fitness of 309.00 with a standard deviation of 48.96 is achieved. The best individual has a fitness of 378.38 which is 94.60% of the maximum achievable fitness. Figure 5 shows how a three-object annular structure is formed in 14190 steps. This structure has a total fitness of 382.70, with a score of 96.94 on the compactness component, 96.55 on the separation component, and 92.27 on the shape component.

At the beginning of a run of the final solution, all of the ants become quickly loaded with an object, and they soon mostly carry central type objects. The reason for this is that the central type objects are only deposited by the ants in neighborhoods that are densely populated with identical objects, and since there are no such neighborhoods at the beginning of an annular sort the ants are unable to get rid of these objects once they pick them up. After some time one or two groups of central type objects begin to emerge. When this happens the ants are not loaded as much of the time as previously and they also begin to carry the other two types of objects. This trend continues as a single cluster of the central object type is formed. When all the central type objects are located in the cluster, the ants begin to put down intermediate type objects along the edges of the cluster. The outermost objects are also moved to some extent at this stage, but they are mostly put down in isolated locations. As more and more of the intermediate type objects are located along the edges of the cluster the outermost objects are moved to a larger extent. Some of these objects are still being put down in isolated locations, but the larger part of the objects are put down along the edges of the cluster. After some time all of the outermost objects are moved to the edges of the cluster and at this time an annular structure has been created.

The formation of an annular structure by the final solution can be summarized as follows. First the ants form a stable cluster of central type objects. Then the ants move intermediate type objects to the edges of the cluster, before doing the same with the outermost objects. The center of the created structure can perhaps be seen as creating an attractive force as objects that become part of the structure are not likely to be removed from the structure. However, some objects are occasionally removed, and the attractive force is thus not absolute. Instead, objects that become

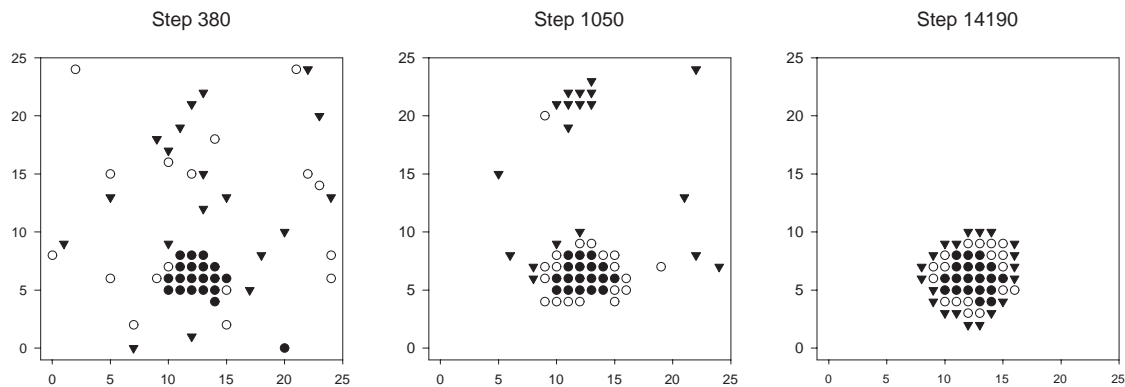


Figure 5: The creation of a three-object annular structure.

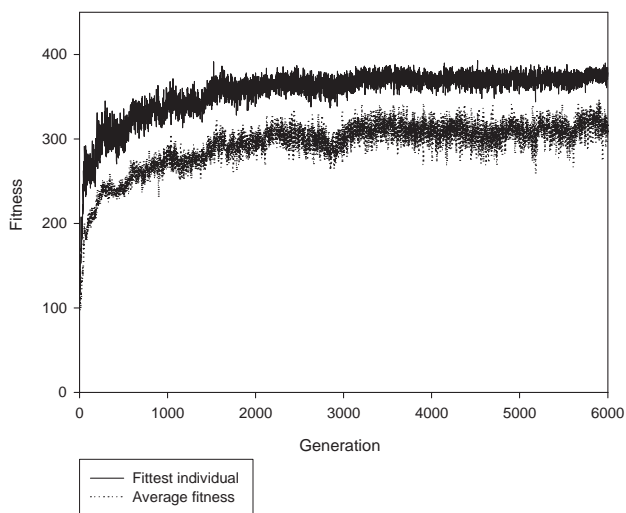


Figure 6: Fitness during the evolution of a three-object annular sort.

part of the structure are unlikely to be removed, and because of this the ants turn to moving the objects that are not yet part of the structure.

When the last solution is tested for stability over 50 different start configurations, that each run for 10000 steps, an average fitness of 329.23 with a standard deviation of 39.09 is achieved.

Figure 6 shows the average fitness of each generation as well as the fitness of the best individual of each generation during the evolution. At the beginning of the evolution (generation 31) the ants continuously form and break up a central type cluster. The ants carry only the central type objects and these are deposited in a single cluster. However, once this cluster is formed the ants begin to break up the cluster and create a new cluster at another location. After the evolution has continued for 13 more generations the ants also move the intermediate type objects. At this stage of the evolution the fittest individual of the generation first creates two or three stable clusters of the central type objects. After these clusters are formed the ants begin to

gather the intermediate type objects along the edges of the already formed clusters. Some of the outermost objects are also eventually moved to the edges of the cluster, but the majority of these objects remain randomly scattered in the environment. After a total of 240 generations the ants of the fittest individual begin by carrying solely central type objects until a single stable cluster of these objects is formed. Then the intermediate and some of the outermost type objects are moved to the edges of this cluster. The objects that are not moved to the edges of the cluster are grouped in small groups containing either one or both types of objects. As the sort continues most of the intermediate type objects are positioned along the edges of the cluster, while groups of outermost type objects are continuously formed and broken down. Many intermediate and outermost type objects are however also deposited by the ants in isolated locations. When the evolution has run for 1420 generations the ants are at the beginning of the sorting no longer carrying solely the central type objects. The ants begin by forming clusters of the central type objects with intermediate type objects along their edges. In addition to these clusters, groups of intermediate type and outermost type objects are also formed. These groups consist either solely of one type, or they consist of one or a few intermediate type objects that are surrounded by outermost type objects. Eventually these groups are broken down and all the objects are at this point part of one or two larger structures. These structures consist of a cluster of central type objects that are surrounded by a band of intermediate objects and then a band of outermost objects. The two surrounding bands are however somewhat noisy, and may contain either type of non-central objects.

5. RESULTS DISCUSSION

There are some interesting similarities between the solutions created for clustering, patch sorting, and annular sorting. In all the solutions there is a phase in the beginning where not much happens. This continues until there is a random fluctuation in the environment in the form of a formation of a small cluster of one, or two, types of objects. This random fluctuation is amplified, and the small cluster grows rapidly. The amplification of random fluctuations is one of the components of self-organizing systems [4].

There is also another important similarity between the patch sorting and annular sorting solutions. During the runs of the solutions the ants focus on one or two of the types

at a time. In the patch sorting the ants group two of the object types before beginning to group the third type. The focus of the ants on only one type of object at a time is even clearer in the annular sorting solutions. In these the ants begin by clustering the central type objects, and when these are clustered the ants begin to pick up objects of the type that should surround this cluster. After this surrounding band is formed the ants will turn to the objects belonging to the next band to be formed. There is thus a clear order in which the ants deposit the items in the growing structures. For the annular sorting solutions the reason for this ordering becomes evident when looking at the different combinations of sensory stimuli that cause the ants to deposit their objects. There are fewer combinations, ie. formations of objects in the ants' eight-cell neighborhoods that cause the ants to deposit the objects on which they focus first. In addition the perceptual inputs that cause the ants to deposit these objects are less likely to be encountered in the beginning of the sort when the objects are mostly scattered. The reason for the order in which the ants focus on objects of different types thus appears to be that the objects on which the ants focus first are harder to 'get rid of' than the objects focused on later in the sorting process. This is also evident in the patch sort solution where the objects the ants focus on first are only deposited in areas where there are other objects of the same type, and at the beginning of the sort there are no such areas.

At the beginning of a sort the ants will pick up objects of any type. However, because of the distribution of objects at this point of the sort the ants are not able to deposit one of the object types. The ants will therefore carry these objects around for a long time until they discover a configuration of objects that allow them to deposit their object. Once the objects become part of the structure the ants will not pick them up again. The ants will then move on to carry the next type of object that is most difficult to deposit. This process continues until the sort is successful. The different sorts thus appears to result from the varying difficulties the ants have with depositing the different types of objects, and the fact that objects that are part of the growing structures are rarely picked up by the ants.

6. SUMMARY AND FUTURE WORK

One of the goals of this work has been to keep things as simple as possible in order to assess the minimum local complexity needed to achieve the global patterns. Because of this the ants all have very limited sensory capabilities. Furthermore, the objects are discrete valued, meaning that the ants can only decide if two objects are identical or different, and not perform any form of similarity measure. The ants are thus capable of performing complex collective tasks with simple individual behaviors using only limited information about the objects in their immediate neighborhood.

A major distinction between this work and previous work in the field is that most previous work has used rules for controlling the behavior of the individual agents. The use of neural networks to control the agents is not a common approach within this field of research, and neither is the use of evolution of swarms of agents. We therefore believe that we are the first to combine the use of neural network controllers and evolution within this field of research.

There are several ways in which our work can be extended in the future. We are currently working on a solution to the

three-object annular sorting problem using a purely local fitness function that resembles the fitness function used to solve the patch-sorting task. So far the solution evolved produces one cluster of the central type and a separate cluster of the two outermost types that is an annular structure. The networks in the solution can also be improved. One obvious improvement would be to reduce the number of input nodes in the networks to always being nine instead of using nine groups of nodes, each containing the same number of nodes as there are different types of objects in the environment. Another improvement would be to create a network where the number of hidden nodes is not dependent on the number of types of objects in the world. A first step towards this is to analyze the function of the hidden nodes in the networks, and we are currently working on this.

7. ACKNOWLEDGEMENTS

We would like to thank Keith Downing for his help during our work and in the writing of this paper.

8. REFERENCES

- [1] G. C. Barker and M. J. Grimson. The physics of muesli. *New Scientist*, 126:37–40, 1990.
- [2] R. Beckers, O. Holland, and J.-L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Proceedings of the Fourth International Conference on Artificial Life*, pages 181–189, 1994.
- [3] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
- [4] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, New Jersey, 2001.
- [5] J.-L. Deneubourg, S. Gross, N. R. Franks, A. Sendova-Franks, C. Detrain, and L. Chrétien. The dynamics of collective sorting: Robot-like ants and ant-like robots. In *Simulation of Adaptive Behavior: From Animals to Animats*, pages 356–363, 1991.
- [6] N. R. Franks and A. Sendova-Franks. Brood sorting by ants: Distributing the workload over the work-surface. *Behavioral Ecology and Sociobiology*, 30:109–123, 1992.
- [7] J. Handl, J. Knowles, and M. Dorigo. Ant-based clustering and topographic mapping. *Artificial Life*, 11(2), 2005.
- [8] O. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5(2):173–202, 1999.
- [9] E. D. Lumer and B. Faieta. Diversity and adaptation in populations of clustering ants. In *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 501–508, 1994.
- [10] C. Melhuish, O. Holland, and S. Hoddell. Collective sorting and segregation in robots with minimal sensing. In *Simulation of Adaptive Behavior: From Animals to Animats 5*, pages 465–470, 1998.
- [11] C. Melhuish, M. Wilson, and A. Sendova-Franks. Patch sorting: Multi-object clustering using minimalist robots. In *Lecture Notes in Artificial Intelligence*, volume 2159, pages 543–552, 2001.
- [12] V. Ramos and J. J. Merelo. Self-organized stigmergic document maps: Environment as mechanism for context learning. In *Proceedings of AEB02 - Spanish Conference on Evolutionary and BioInspired Algorithms*, pages 284–293, 2002.
- [13] A. Sendova-Franks, S. R. Scholes, N. R. Franks, and C. Melhuish. Brood sorting by ants: Two phases and differential diffusion. *Animal Behaviour*, 68(5):1095–1106, 2004.
- [14] M. Wilson, C. Melhuish, A. Sendova-Franks, and S. R. Scholes. Algorithms for building annular structures with minimalist robots inspired by brood sorting in ant colonies. *Autonomous Robots*, 17:115–136, 2004.