

# Complexity-Seeking Ants

Howard Gutowitz

ESPCI; Laboratoire d'Electronique  
10 rue Vauquelin; 75005 Paris, France  
gutowitz@amoco.saclay.cea.fr

## Abstract

Deneubourg et al. (1991) introduced a model of clustering behavior in ants. They found that simple model ants were able gather into piles objects initially strewn randomly across the plane. The model is in qualitative agreement with the behavior of real ants. The model ants operate according to local strategic rules and possess only local perceptual capacities. Nonetheless, they are able to impose global order. What is the mechanism underlying this phenomenon? We hypothesize that it is the combination of two processes, one which decreases small-scale complexity, and another which couples small-scale decreases in complexity to larger scales. To test this hypothesis, we introduce variant ants which have a complexity-seeking strategy. These ants can "see" local complexity, and tend to perform actions (picking up and putting down objects) in regions of highest local complexity. Using this strategy, they are able to accomplish their task more efficiently than Deneubourg et al.'s basic ants. For both basic and complexity-seeking ants, we find that complexity-reduction begins at the finest scale and propagates out to ever-increasing scales. Reduction of spatial entropy is used as a global measure of clustering success. This global measure correlates well with a local measure of success which the individual ants themselves can calculate: the proportion of object depositions they perform which decrease the local complexity.

### Introduction

Organisms can collect energy and focus it on the reduction of entropy in their environment. We study this organizing capability of living things in the context of a simple model of collective behavior in ants introduced by Deneubourg et al. (1991). The ants of Deneubourg et al. are able to impose order in their environment using most rudimentary perceptual, motor, and strategic mechanisms. We will examine a number of variants to Deneubourg et al.'s basic ants. We shall see that some of these variants succeed better at their task than the basic ants do. However, our ultimate goal lies beyond the engineering of better ants, in the direction of connecting local and global criteria for the evaluation of autonomous-agent strategies (for some related work see Meyer and Wilson, (1991) and Huberman (1988)).

The global measure we will use to evaluate the success of clustering strategies is a standard statistical-mechanical measure: the coarse-grained spatial entropy of the environment. As the ants can only perform local computations, they themselves cannot calculate global spatial entropy. As individuals they are unable to know how much contribution they make to the collective task. We will therefore build a local measure of success from quantities the ants themselves can perceive, quantities related to their own internal functions, and local effects they can make in their environment. Though we will not directly consider learning in ants, it is clear that some such local measure will be required as the basis of any learning algorithm which is to operate at the individual, rather than the population level.

### The Basic Ants

We will refer to the model ants introduced by Deneubourg et al. as basic ants. Basic ants have 1) a finite memory, 2) an object-manipulation capacity, the capacity pick up or put down objects, 3) a function which gives the probability to manipulate an object as a function of the values in memory and a random variable, and 4) the capacity to execute a brownian motion. The memory is simply a register of length  $n$  in which is recorded the presence or absence of objects at the ant's previous  $n$  locations. Throughout the present work  $n$  is set to 15. At each time step a basic ant generates a random number  $p$  between 0 and 1. It manipulates objects as a function of  $p$  and a threshold calculated according to

$$P(\text{pick up}) = \frac{k^+}{(k^+ + f)^2}; P(\text{put down}) = \frac{f}{(k^- + f)^2}, \quad (1)$$

where  $f$  is the fraction of memory locations registering the presence of an object (an estimate of the local density of objects), and  $k^+$  and  $k^-$  are parameters, set by Deneubourg et al. to 0.3 and 0.1 respectively. In the present work the values  $k^+=0.43, k^-=0.12$  are used for complexity-seeking ants, and  $k^+=0.38, k^-=0.15$  for the basic ants, as these values were found by simulation to be near-optimal.  $P(\textit{put down})$ ,  $P(\textit{pick up})$  are computed if the ant does, respectively does not currently hold an object. If  $p$  exceeds  $P$  then the action is taken, otherwise it is not. Basic ants (as well as the other ants considered here) do not put down objects at already occupied sites.

In the work of Deneubourg et al. the basic ants are used to collect objects strewn on a square grid into piles. A close variant of the basic ants is used to sort two different types of objects into piles of objects of a single type. In this paper we will only consider the collection of objects into piles without respect to type. We refer to this process as clustering, with the understanding that the same approach could be taken to the full sorting problem.

The speed of clustering depends in general on the ratio of the number of ants to the number of objects to be clustered. In the limit where this ratio is small, one may expect a linear relationship between the number of ants and the speed of clustering. When the ratio is high, i.e. when at any given time a significant fraction of the objects are being carried by ants, this linear relationship will break down. Here, as in the work of Deneubourg et al., we study the limit of small ( $\#ants$ )/( $\#objects$ ) ratio. We use 25 ants for approximately  $3.3 \times 10^4$  objects.

#### Perception of Local Complexity

The ability of basic ants to reduce global environmental entropy can be traced to their primitive capacity to perceive complexity locally and behave accordingly. Complexity and entropy are closely related concepts. While entropy can be reliably measured on a global scale, it is difficult to find an adequate measure of entropy on a local scale. In this section we assert that the basic ants of Deneubourg et al. already have some implicit capacity to measure local entropy, and introduce a function, which we will call local complexity, which is more explicitly related to local entropy.

Basic ants record in memory the presence or absence of an object at the last  $n$  lattice sites they have visited. From this record they calculate the local density of objects on the lattice. Entropy is related to the density. At either very high or very low density the entropy is low, while at intermediate densities the entropy is high. Basic ants pick up or put down objects in a

probabilistic fashion according to density thresholds. If the density is very low or high, then the ant is likely to perform the correct action, picking up in regions of low density and putting down in regions of high density. If the estimated local density is close to a threshold, then the certitude of performing the appropriate action (picking up or putting down) is low.

We build upon this insight by engineering ants which have an explicit capacity to sense complexity. Given a local configuration of states, complexity can be estimated in various fashions. We choose a simple method based on a 9-cell neighborhood about the ant's current location. A 9-cell neighborhood in the square lattice has 12 interior faces between cells. The "complexity" of the neighborhood is taken as the number of faces which separate cells of different type—containing or not containing an object. Thus, either an "all-empty" or "all-occupied" neighborhood will have complexity 0, while a checkerboard pattern will have complexity 12. Complexity-seeking ants can calculate the complexity of their current position, thus they can "see" the ambient complexity of their local environment.

Complexity-seeking ants appeal to their sensory system to direct both their physical motion and their object-manipulation activities. At each step of simulation, each ant calculates its local complexity,  $C$ . The information in the value  $C$  can be exploited either deterministically or probabilistically. In a deterministic strategy, the value  $C$  is compared with a threshold. If  $C$  exceeds the threshold, then the ant randomly chooses a new direction of motion and decides whether to pick up or put down an object at its current position. In a probabilistic strategy, the ant generates a random number between 0 and 12. If the number is greater than  $12 - (C + T)$ , where  $T$  is a threshold, then the ant will randomly choose a new direction of motion and execute the object manipulation strategy used by the basic ants. If the number is less than  $12 - (C + T)$ , the ant will continue along its previous direction of motion, and will not attempt to manipulate an object. The result of this rule is that complexity-seeking ants tend to manipulate objects only in regions of high complexity. Similarly, the complexity-seeking ants execute the random-move generator of the basic ants only when the complexity is high. In regions of low complexity, such as the large spaces between piles in mature environments, the ants generally continue in the direction they were previously moving. The result is that complexity-seeking ants spend less time in futile search in the spaces between piles than their basic-ant counterparts, and spend correspondingly more time in careful processing at the borders of

piles.

In these simulations, complexity-sensing affects both the movement and object-manipulation strategies of the ants. No attempt is made here to distinguish between these effects. This will be taken up in a subsequent report.

We will consider a number of different threshold values for both deterministic and probabilistic complexity-seeking strategies. For deterministic strategies  $T$  values of 0, 3 and 6 are studied, while for the probabilistic strategies  $T$  values of 0 and 0.5 are studied.

The patterns resulting from running each of these strategies  $10^7$  time steps on a random initial pattern with density 0.1 is shown in figure 1. Note that periodic boundary conditions are used (the ants are considered to live on a torus). The panels are arranged top to bottom and left to right in order of increasing determinism in the strategies.

After  $10^7$  time steps, the basic ants (figure 1 top left) have transformed the uniformly random initial pattern into a mottled pattern of many small, loose piles. The probabilistic complexity-seeking ants,  $T = 0.5$  (figure 1 middle left) have in the same time produced a pattern with one nearly continuous, loose pile. The probabilistic complexity-seeking ants,  $T = 0$  (figure 1 bottom left) have collected the objects into two, tight piles. It would seem at this point that this strategy is by far the best, however see next section. The second column of figure 1 shows the results of the deterministic complexity seeking strategies with  $T = 0, 3, 6$ , in top, middle, bottom, positions respectively. While the deterministic strategy with  $T = 0$  produces a good initial clustering, the  $T = 3$  and  $T = 6$  strategies are grossly inefficient.

The rectangular piles produced by the deterministic complexity-seeking  $T = 3$  ants are artifacts of the complexity measure we have adopted. When the threshold is 3 and the rule is deterministic, ants can only put down or pick up objects from the corner of a rectangular pile. Thus rectangular piles are stable against the activity of the ants. The creation of such artifacts in the case of highly deterministic strategies suggests that there is a lower bound for the rate of random input required for good clustering. While the introduction of a local complexity measurement has reduced the random-input requirements of the basic ants, it has not eliminated it. It is evident that to *any* fully deterministic strategy there corresponds a set of configurations which cannot be clustered.

### Spatial Entropy

To quantitatively study the clustering phenomenon shown in figure 1, we






Figure 1: Patterns produced by some variant ant strategies. All patterns produced after  $10^7$  time steps on a 256 square grid. From top to bottom, column 1: Basic ants, probabilistic complexity-seeking  $T = 0.5$  and  $T = 0$ . column 2: deterministic complexity-seeking  $T = 0, 3, 6$ . Periodic boundary conditions.

compute the evolution with time of the coarse-grained entropy of the patterns. The coarse-grained spatial entropy is taken as  $\sum_{grains} P_{grain} \text{Log}(P_{grain})$ , where  $P_{grain}$  is the fraction of the objects on the lattice which are found in the given grain ( $\sum_{grains} P_{grain} = 1$ ). The entropy is measured at intervals during the clustering activity of the ants. Several different levels of coarse graining are used in order to track the organization effected by the ants across length scales. In typical simulations a 256x256 square 2-D lattice is used, with periodic boundary conditions imposed. 5 levels of coarse graining are computed, varying from a 8x8- to 128x128-site grains. To control inaccuracies due to clusters which span grain boundaries, the location of the entropy-computation grid is shifted relative to the lattice, and a minimum over shifts of the entropy is taken.

We see that in all cases entropy measured at the smallest grain sizes decreases quickly to an equilibrium value. The smaller the grain size, the faster the approach to equilibrium. This indicates that organization begins at the smallest scales and then propagates to larger scales. Stochastic fluctuations due to probabilistic strategy elements facilitate the diffusion of organization




Figure 2: Evolution of the Coarse-Grained Entropy. Top: basic ants. Bottom: probabilistic complexity-seeking ants,  $T = 0.5$ . The line thickness increases as a function of grain size. Note difference in vertical scale between panels.

from small to large spatial scales.

In figure 2 the evolution of the coarse-grained entropy at all space scales is shown for basic and probabilistic complexity-seeking  $T = 0.5$  ants. For both basic and probabilistic complexity-seeking ants, the reduction in entropy effected by the ants is seen first in the small scales and later at larger scales. Small-scale entropy quickly reaches an equilibrium value while the large-scale entropy continues to gradually decrease. The complexity-seeking ants achieve lower values of equilibrium spatial entropy, at all scales, than the basic ants. The reduction in entropy occasionally proceeds by jumps. This is especially clear in the evolution of the coarsest-grained entropy under the action of the complexity-seeking ants. These jumps follow plateaus during which similarly sized piles compete for deposition of objects by the ants. In every case, stochastic fluctuations will eventually lead some of the piles to sufficiently dominate the others in size so that the ants act to amplify the difference by positive feedback. This destroys the smaller piles and leads to a new metastable state during which competition occurs at a larger space scale.

We now examine the evolution of the coarse-grained entropy for four of the variant strategies considered in this paper. In figure 3, the entropy at the largest scale of coarse-graining is shown for each of these variant strategies. This figure should be compared to figure 1.

The probabilistic strategy with  $T = 0$  is initially a much better clusterer than the other strategies. The probabilistic strategy with  $T = 0.5$  wins over the deterministic strategy, but less impressively. However, in the end game as the environment matures, the situation becomes more complicated. When the number of piles has been reduced to 2, which occurs between 10 and 30 million time steps for all complexity-seeking strategies in this figure, the more deterministic strategies cannot cope well with the task of choosing which pile to favor with object deposition. The probabilistic  $T = 0.5$  strategy also enters a meta-stable state consisting of two piles, as we saw in figure 2, but it is able to gradually break the symmetry between the piles. When  $T$  is set to 0, the coarse-grained entropy may actually significantly increase for long periods, as happens in this run at approximately  $2.4 \times 10^7$  time steps into the simulation.

Part of the reason for the failure of the more deterministic strategies can be seen by considering the action of the ants in a mature environment with large spaces between piles. For the deterministic strategy, and the probabilistic strategy with  $T = 0$ , no changes of direction occur in the empty spaces between piles. Ants thus continue in a straight line through regions between piles. If there are few piles, some ants will be trapped into orbits which never intersect a pile. These ants perform no work. This "compulsive" behavior which a detriment at large time has positive value at short time. A population of compulsive ants produces tighter piles than a population of their well-adapted counterparts, and do so in a comparatively short time (see figure 1).

#### Strategic Efficacy

We considered a number of different local measures of ant strategies. Two of these will be described here. We give the ants the ability to measure the local complexity of a site both before and after it puts down an object, as well as an integer counter. Thus it can track its rate of success in reducing the local complexity when it puts down an object. If an ant puts down an object and thereby decreases the local complexity, it is considered to have scored a success. The rate of success computed over intervals of 20,000 time steps is shown as a function of time in figure 4 for four clustering strategies.






Figure 3: Evolution of the Coarse-grained Entropy for the Variant Strategies. This figure shows the evolution of the coarse-grained entropy for each of the variant strategies: basic, probabilistic complexity-seeking  $T = 0$  and  $T = 0.5$ , and deterministic complexity-seeking with  $T = 0$ .

The basic ants and each flavor of complexity-seeking ants were run  $10^7$  time steps on a  $256 \times 256$  grid. Remarkably, the raw number of successes in an interval indicates little concerning the success of the strategy evaluated at a global level. Indeed, the strategy which has the greatest initial success globally, the probabilistic strategy with threshold 0, scores the lowest rate of success measured locally.

If the rate of success is compared with the raw rate of putting down objects, a picture consistent with the global measure of success emerges. This is shown in figure 4. Here we see that the probabilistic  $T = 0$  ants rapidly reach an equilibrium rate of success per put down of nearly 40 percent. Recall that these are the best clusterers initially as judged by the rate of decrease of the spatial entropy of the lattice. Probabilistic ants  $T = 0.5$  eventually reach the same rate of success, but take longer to do so. This again correlates with the global success measure. The deterministic  $T = 0$  strategy initially has a local rate of success per object deposition similar to but somewhat below the probabilistic  $T = 0.5$  ants. The deterministic strategy appears to reach an asymptotic rate of local success which is lower than the probabilistic




Figure 4: Local Measures of Success. a) The raw rate of success achieved by each of 4 sorting strategies. b) The ratio between the rate of success and the rate of putting down objects.

strategies. Finally, the basic ants have a very low rate of success per put-down compared to the complexity-seeking ants. This too correlates well with the global measure of success.

### Discussion

The clustering behavior discovered by Deneubourg et al. is "emergent" in the sense that ants are not explicitly programmed to form clusters. Clustering results from the dependence of the ant's behavior on the perceived local density. In much the same way, the complexity-seeking behavior described here is emergent from the dependence of the ant's behavior on the perceived local complexity. As with clustering itself, the complexity-seeking behavior inherits robustness to perturbations from the implicit means by which it is programmed. Clustering is representative of a large class of biological activities which lead to structure in an environment. Appropriate adaptation of the complexity-seeking strategy could lead to new methods from improving the efficiency of many types of collective computation. To prepare for this, one should study how clustering depends on the topology of the lattice, the number of types of objects to be clustered and sorted, their relative concentrations etc.

Our experimental results show that the global success measure, reduction of spatial entropy, correlates well with a local measure of success, the proportion of object deposition which reduce local complexity. The rate of successful

object deposition by itself does not correlate with the global measure. It appears that the best strategy in this game is to invest increasingly more time and energy into achieving each success as the environment matures. Churning the success counter by inconsiderately picking up and putting down objects does not lead to effective global behavior. While complexity-seeking increases the effectiveness of clustering and reduce the amount of random input required, there appears to be a lower bound to the amount of random input needed to adapt to randomness in the environment.

The complexity-seeking model explored in this paper could be compared with biological experiments. It makes predictions concerning the statistics of activity of real ants as they go about forming clusters. For instance, it may be possible to obtain laboratory measurements of the amount of time real ants spend in various environments as a function of local complexity. A valuable measure is the probability for an ant to change direction of motion as a function of the local environmental complexity. While it is unlikely that real ants compute the particular, ad-hoc, local complexity function employed in these simulations, it is reasonable that they would have some means to evaluate the general simplicity or complexity of their immediate sensory environment.

The clustering problem is a good means to connect issues in computational sociology and ethology with the tools of statistical mechanics and dynamical systems theory. On one hand it presents entities whose activities have a clear biological interpretation. On the other hand, it exhibits phenomena, such as a type of order/disorder transition, which are familiar from physics. One hopes that this connection will help organize and direct the study of autonomous agents.

## References

- [1] Deneubourg, J.L., Goss, S., Franks, N., Sendova-Franks A., Detrain C., and Chrétien, L. 1991. The Dynamics of Collective Sorting, Robot-Like Ants and Ant-Like Robots pg 356 in [2].
- [2] J.-A. Meyer, J.-A., and Wilson, S.W., 1991. Eds. From Animals to Animats. MIT Press, Bradford Books.
- [3] Huberman. B.A. Ed. 1988. The Ecology of Computation, North-Holland, Amsterdam.