



Genetic Programming and Evolvable Machines, 6, 111–125, 2005
© 2004 Kluwer Academic Publishers. Manufactured in The Netherlands.

1 Emergence of Collective Behavior in Evolving 2 Populations of Flying Agents

3 LEE SPECTOR

4 *School of Cognitive Science, Hampshire College, Amherst, MA 01002, USA*

5 JON KLEIN

6 *School of Cognitive Science, Hampshire College, Amherst, MA 01002, USA; Physical Resource Theory, Chalmers*

7 *University of Technology and Göteborg University, SE-412 96, Göteborg, Sweden*

8 CHRIS PERRY

9 MARK FEINSTEIN

10 *School of Cognitive Science, Hampshire College, Amherst, MA 01002, USA*

11 *Submitted ; Revised*

Au: Please
provide
dates.

12 **Abstract.** We demonstrate the emergence of collective behavior in two evolutionary computation systems, one
13 an evolutionary extension of a classic (highly constrained) flocking algorithm and the other a relatively un-
14 constrained system in which the behavior of agents is governed by evolved computer programs. The first system
15 demonstrates the evolution of a form of multicellular organization, while the second demonstrates the evolution
16 of a form of altruistic food sharing. In this article we describe both systems in detail, document the emergence
17 of collective behavior, and argue that these systems present new opportunities for the study of group dynamics
18 in an evolutionary context. We also provide a brief overview of the BREVE simulation environment in which the
19 systems were produced, and of BREVE's facilities for the rapid, exploratory development of visualization strategies
20 for artificial life.

21 **Keywords:** collective behavior, 3D virtual worlds, visualization

22 1. Evolution of collective behavior

23 The evolution of group behavior is a central concern in evolutionary biology and behavioral
24 ecology. Ethologists have articulated many costs and benefits of group living and have
25 attempted to understand the ways in which these factors interact in the context of evolving
26 populations. For example, they have considered the thermal advantages that warm-blooded
27 animals accrue by being close together, the hydrodynamic advantages for fish swimming
28 in schools, the risk of increased incidence of disease in crowds, the risk of cuckoldry
29 by neighbors, and many advantages and risks of group foraging [5]. Attempts have been
30 made to understand the evolution of group behavior as an optimization process operating
31 on these factors, and to understand the circumstances in which the resulting optima are
32 stable or unstable [7, 13]. Similar questions arise at a smaller scale and at an earlier phase
33 of evolutionary history with respect to the evolution of symbiosis, multicellularity, and
34 other forms of aggregation that were required to produce the first large, complex life forms
35 [1, 6].

Game theoretic simulations, often based on the Prisoner's Dilemma, have provided ample data and insights concerning the evolution of group behavior, although usually at a level of abstraction far removed from the physical risks and opportunities presented by real environments (see, e.g., [2], about which we say a bit more below). One line of game-theoretic research that is related to the results presented here explores the minimal requirements for the evolution of altruism. Of particular note, Riolo et al. showed that altruism can sometimes evolve even when agents have no knowledge of the relatedness, reputation, or propensity for reciprocation of other agents; the agents in their system based all decisions to cooperate or defect solely upon the differences between the arbitrary numeric "tags" attached to themselves and to other agents [12].¹

Artificial life technologies provide new tools for the investigation of these issues. One well-known, early example was the use of the Tierra system to study the evolution of a simple form of parasitism [8]. Other investigators have attempted to study the evolution of collective behavior in populations of flying or swimming agents that are similar in some ways to those investigated here, with varying degrees of success [10, 16]. The latest wave of artificial life technology presents yet newer opportunities, however, as it is now possible to conduct much more elaborate simulations on modest hardware and in short time spans, to observe both evolution and behavior in real time in high-resolution 3D displays, and to interactively explore the ecology of evolving ecosystems.

In this article we describe two recent experiments in which the emergence of collective behavior was observed in evolving populations of flying agents. The first experiment used a system, called SWARMEVOLVE 1.0, that extends a classic flocking algorithm to allow for multiple species, goal orientation, and evolution of the constants in the hard-coded motion control equation. In this system we observed the emergence of a form of collective behavior in which species act similarly to multicellular organisms. The second experiment used a later and much-altered version of this system, called SWARMEVOLVE 2.0, in which the behavior of agents is controlled by evolved computer programs instead of a hard-coded motion control equation.² In this system we observed the emergence of altruistic food-sharing behaviors and investigated the link between this behavior and the stability of the environment.

Both SWARMEVOLVE 1.0 and SWARMEVOLVE 2.0 were developed within BREVE, a simulation package designed by Klein for realistic simulations of decentralized systems and artificial life in 3D worlds [4]. In the next section we provide a brief overview of BREVE and of the facilities that it provides for development and visualization of artificial life simulations. In the subsequent sections we describe the two SWARMEVOLVE systems and the collective behavior phenomena that we observed within them. This is followed by some brief remarks about the potential for future investigations into the evolution of collective behavior using artificial life technology.

2. BREVE

BREVE is an open-source simulation environment³ that simplifies the construction of advanced artificial life and multi-agent simulations. Users define individual agent behaviors in a simple object-oriented language and then explore the behaviors of the agents by means of a rich visualization engine. All aspects of the simulation, including object and memory

79 management, communication between agents, and integration are automatically handled by
80 the BREVE engine. BREVE simplifies the programming of advanced agent behaviors through
81 integrated support for matrix and vector arithmetic, collision and “neighbor” detection, re-
82 alistic physical simulation, and facilities for customization and expandability via a dynamic
83 plugin architecture.

84 BREVE promotes rapid programming and prototyping of simulations largely through the
85 use of a built-in, high-level, object-oriented programming language called STEVE. Some
86 simulation environments, such as SFI Swarm⁴ and RePast,⁵ use existing languages such as
87 Java and Objective-C to define agent behaviors. By incorporating an integrated programming
88 language designed for 3D simulation, BREVE avoids much of the overhead associated with
89 implementing simulations in these environments. In addition, the integrated approach allows
90 new features to be built directly into the syntax of the language, instead of requiring users
91 to learn a new programming API. For example, 3D vectors, 3D matrices and hash tables
92 are all included in BREVE as native types and are accessed directly as part of the language.
93 BREVE’s strict object-oriented design framework also enables rapid simulation development
94 by encouraging code reuse and modularity. The demonstration simulations included with
95 the standard BREVE distribution provide simple examples of a wide variety of simulation
96 paradigms (including real 3D physical simulation, 2D and 3D cellular automata, particle
97 systems, flocking simulations and Braitenberg vehicle simulations). Code from these demos
98 can be readily reused and modified, allowing rapid simulation development based on existing
99 examples.

100 While simulation environments such as Swarm and StarLogo [9] also facilitate multi-
101 agent simulation development, BREVE arguably offers more sophisticated options for visu-
102 alization and spatial simulations. In terms of spatial simulations, Swarm and StarLogo are
103 generally restricted to simulations taking place on discrete 2D grids, which tend to be most
104 useful in dealing with 2D cellular automata and simple 2D multi-agent simulations. While
105 both environments also offer additional visualization tools such as graphs and charts, these
106 features are generally used to display statistics and information about a simulation, and
107 not to render a representation of the actual simulation space. BREVE, on the other hand, is
108 suitable for continuous and discrete simulations in both 2D and 3D, as well as simulations
109 which combine aspects of both. Although BREVE represents all simulations internally using
110 a 3D continuous space, discrete simulations can be built by employing grids of “patches,”
111 which are rectangular blocks in 3D space. These grids can have any dimensions, can be po-
112 sitioned in 3D space like any other BREVE object, and can interact freely with regular objects
113 or even other patch grids. In conjunction with agents moving through continuous space, the
114 patch grids also allow for partitioning of 3D space into different regions representing, for
115 example, different temperatures or concentrations of chemicals.

116 BREVE’s OpenGL visualization engine portrays the locations and behaviors of agents in
117 real-time. Visualization of spatial simulations allow observers to readily discover patterns
118 that might easily be overlooked with simple numerical output. The visualization engine is
119 thus critical with respect to exploratory artificial life simulation development.

120 In addition to the “literal” spatial visualization features which portray position and mo-
121 tion of objects in 3D space, BREVE also provides a variety of “special effects.” These special
122 effects are not only cosmetic, but can also be used to convey high-dimensional data about
123 the state of agents in a simulation through visual means. Integrated special effects in-
124 clude shadows, reflections, fog, textures, displaying objects as semi-transparent bitmaps,

displaying objects as light sources (“lightmaps”), and drawing lines connecting arbitrary objects. Manipulation of the color and transparency of objects in the simulation, for example, may be used to display dynamic information about an agent’s behavior or fitness, and shadows and reflections can be used to add visual distance cues. All of the features of the visualization engine, including lighting, special effects and camera position, can be controlled dynamically. These features can be manipulated manually by the user in order to gain a better perspective on the simulation, or programmatically by the simulation itself (via STEVE code). The SWARMEVOLVE systems described below employ automatic control of several of these features to provide an information-rich visual interface that displays, for example, velocity-based agent orientation, energy-based agent color/size, death rates via piles of “corpses,” and close-up views of the center of activity via automatically adjusted camera zoom and targeting. In additional experiments we explored the system from an “agent’s eye view” by placing the camera in an agent, and “sharing networks” by drawing lines dynamically between agents that share energy (as in Figure 6, below). The descriptions of BREVE and STEVE in this article are necessarily incomplete; please see [4] for further details.⁶

3. SWARMEVOLVE 1.0

One of the demonstration programs distributed with BREVE is SWARM, a simulation of flocking behavior modeled on the “boids” work of Craig W. Reynolds [11]. In the BREVE SWARM program the acceleration vector for each agent is determined at each time step via the following formulae:

$$\mathbf{V} = c_1\mathbf{V}_1 + c_2\mathbf{V}_2 + c_3\mathbf{V}_3 + c_4\mathbf{V}_4 + c_5\mathbf{V}_5$$

$$\mathbf{A} = m \left(\frac{\mathbf{V}}{|\mathbf{V}|} \right)$$

The c_i are positive constants and the \mathbf{V}_i are vectors determined from the state of the world (or in one case from the random number generator) and then normalized to length 1. \mathbf{V}_1 is a vector away from neighbors that are within a “crowding” radius, \mathbf{V}_2 is a vector toward the center of the world, \mathbf{V}_3 is the average of the agent’s neighbors’ velocity vectors, \mathbf{V}_4 is a vector toward the center of gravity of all agents, and \mathbf{V}_5 is a random vector. In the second formula we normalize the resulting velocity vector to length 1 (assuming its length is not zero) and set the agent’s acceleration to the product of this result and m , a constant that determines the agent’s maximum acceleration.⁷ The system also models a floor and hard-coded “land” and “take off” behaviors, but these are peripheral to the focus of this article. By using different values for the c_i and m constants (along with the “crowding” distance, the number of agents, and other parameters) one can obtain a range of different flocking behaviors; many researchers have explored the space of these behaviors since Reynolds’s pioneering work [11].

SWARMEVOLVE 1.0 enhances the basic BREVE SWARM system in several ways. First, we created three distinct species⁸ of agents, each designated by a different color. As part of this enhancement we added a new term, $c_6\mathbf{V}_6$, to the motion formula, where \mathbf{V}_6 is a vector away from neighbors of *other species* that are within a “crowding” radius. Goal-orientation was

163 introduced by adding a number of randomly moving “energy” sources to the environment
164 and imposing energy dynamics. As part of this enhancement we added one more new term,
165 $c_7 \mathbf{V}_7$, to the motion formula, where \mathbf{V}_7 is a vector toward the nearest energy source. Each
166 time an agent collides with an energy source it receives an energy boost (up to a maximum),
167 while each of the following bears an energy cost:

- 168 – Survival for a simulation time step (a small “cost of living”).
- 169 – Collision with another agent.
- 170 – Being in a neighborhood (bounded by a pre-set radius) in which representatives of the
- 171 agent’s species are outnumbered by representatives of other species.
- 172 – Giving birth (see below).

173 The numerical values for the energy costs and other parameters can be adjusted arbitrarily
174 and the effects of these adjustments can be observed visually and/or via statistics printed
175 to the log file; values typical of those that we used can be found in the source code for
176 SWARMEVOLVE 1.0.⁹

177 As a final enhancement we leveraged the energy dynamics to provide a fitness function
178 and used a genetic encoding of the control constants to allow for evolution. Each individual
179 has its own set of c_i constants; this set of constants controls the agent’s behavior (via the
180 enhanced motion formula) and also serves as the agent’s genotype. When an agent’s energy
181 falls to zero the agent “dies” and is “reborn” (in the same location) by receiving a new
182 genotype and an infusion of energy. The genotype is taken, with possible mutation (small
183 perturbation of each constant) from the “best” current individual of the agent’s species
184 (which may be at a distant location).¹⁰ We define “best” here as the *product* of energy and
185 age (in simulation time steps). The genotype of the “dead” agent is lost, and the agent that
186 provided the genotype for the new agent pays a small energy penalty for giving birth. Note
187 that reproduction is asexual in this system (although it may be sexual in SWARMEVOLVE
188 2.0).

189 The visualization system presents a 3D view (automatically scaled and targeted) of the
190 geometry of the world and all of the agents in real time. Commonly available hardware is
191 sufficient for fluid action and animation. Each agent is a cone with a pentagonal base and a
192 hue determined by the agent’s species (red, blue, or purple). The color of an agent is dimmed
193 in inverse proportion to its energy — agents with nearly maximal energy glow brightly while
194 those with nearly zero energy are almost black. “Rebirth” events are visible as agents flash
195 from black to bright colors.¹¹ Agent cones are oriented to point in the direction of their
196 velocity vectors. This often produces an appearance akin to swimming or to “swooping”
197 birds, particularly when agents are moving quickly. Energy sources are flat, bright yellow
198 pentagonal disks that hover at a fixed distance above the floor and occasionally glide to
199 new, random positions within a fixed distance from the center of the world. An automatic
200 camera control algorithm adjusts camera zoom and targeting continuously in an attempt to
201 keep most of the action in view.

202 Figure 1 shows a snapshot of a typical view of the SWARMEVOLVE world. An animation
203 showing a typical action sequence can be found on-line.¹²

204 SWARMEVOLVE 1.0 is simple in many respects but it nonetheless exhibits rich evolution-
205 ary behavior. One can often observe the species adopting different strategies; for example,
206 one species often evolves to be better at tracking quickly moving energy sources, while



Figure 1. A view of SWARMEVOLVE 1.0 (which is in color but will print black and white in the journal). The agents in control of the pentagonal energy source are of the purple species, those in the distance in the upper center of the image are blue, and a few strays (including those on the left of the image) are red. All agents are the same size, so relative size on screen indicates distance from the camera.

another evolves to be better at capturing static energy sources from other species. An animation demonstrating evolved strategies such as these can be found on-line.¹³

4. Emergence of collective behavior in SWARMEVOLVE 1.0

Many SWARMEVOLVE runs produce at least some species that tend to form static clouds around energy sources. In such a species, a small number of individuals will typically hover within the energy source, feeding continuously, while all of the other individuals will hover in a spherical area surrounding the energy source, maintaining approximately equal distances between themselves and their neighbors. Figure 2 shows a snapshot of such a situation, as does the animation at <http://hampshire.edu/lspector/swarめevolve-ex2.mov>; note the behavior of the purple agents.

We initially found this behavior puzzling as the individuals that are not actually feeding quickly die. On first glance this does not appear to be adaptive behavior, and yet this behavior emerges frequently and appears to be relatively stable. Upon reflection, however, it was clear that we were actually observing the emergence of a higher level of organization.

When an agent dies it is reborn, in place, with a (possibly mutated) version of the genotype of the “best” current individual of the agent’s species, where quality is determined from the product of age and energy. This means that the new children that replace the dying individuals on the periphery of the cloud will be near-clones of the feeding individuals within the energy source. Since the cloud generally serves to repel members of other species, the formation of



Figure 2. A view of SWARMEVOLVE 1.0 in which a cloud of agents (the blue species) is hovering around the energy source on the right. Only the central agents are feeding; the others are continually dying and being reborn. As described in the text this can be viewed as a form of emergent collective organization or multicellularity. In this image the agents controlling the energy source on the left are red and most of those between the energy sources and on the floor are purple.

226 a cloud is a good strategy for keeping control of the energy source. In addition, by remaining
 227 sufficiently spread out, the species limits the possibility of collisions between its members
 228 (which have energy costs). The high level of genetic redundancy in the cloud is also adaptive
 229 insofar as it increases the chances that the genotype will survive after a disruption (which
 230 will occur, for example, when the energy source moves).

231 The entire feeding cloud can therefore be thought of as a genetically coupled collective,
 232 or even as a multicellular organism in which the peripheral agents act as defensive organs
 233 and the central agents act as digestive and reproductive organs.

234 5. SWARMEVOLVE 2.0

235 Although SWARMEVOLVE 2.0 was derived from SWARMEVOLVE 1.0 and is superficially
 236 similar in appearance, it is really a fundamentally different system.

237 The energy sources in SWARMEVOLVE 2.0 are spheres that are depleted (and shrink) when
 238 eaten; they re-grow their energy over time, and their signals (sensed by agents) depend on
 239 their energy content and decay over distance according to an inverse square law. Births occur
 240 near mothers and dead agents leave corpses that fall to the ground and decompose; a view
 241 of such corpses is shown in Figure 3. A form of energy conservation is maintained, with
 242 energy entering the system only through the growth of the energy sources. All agent actions
 243 are either energy neutral or energy consuming, and the initial energy allotment of a child
 244 is taken from the mother. Agents get “fatter” (the sizes of their bases increase) when they



Figure 3. A pile of corpses forming under a dangerous area in SWARMEVOLVE 2.0.

have more energy, although their lengths remain constant so that length still provides the appropriate cues for relative distance judgement in the visual display; Figure 4 shows some of these effects. A graphical user interface has also been added to facilitate the experimental manipulation of system parameters and monitoring of system behavior.

The most significant change, however, was the elimination of hard-coded species distinctions and the elimination of the hard-coded motion control formula (within which, in SWARMEVOLVE 1.0, only the constants were subject to variation and evolution). In SWARMEVOLVE 2.0 each agent contains a computer program that is executed at each time step. This program produces two values that control the activity of the agent:

1. a vector that determines the agent's acceleration,
2. a floating-point number that determines the agent's color.

Agent programs are expressed in Push, a programming language designed by Spector to support the evolution of programs that manipulate multiple data types, including code; the explicit manipulation of code supports the evolution of modules and control structures, while also simplifying the evolution of agents that produce their own offspring rather than relying on the automatic application of hand-coded crossover and mutation operators [14, 15].

The Push instructions available for use in agent programs are shown in Table 1. In addition to the standard Push instructions, operating on integers, floating point numbers, Boolean values, and code expressions, instructions were added for the manipulation of vectors and for SWARMEVOLVE-specific sensors and actions. Note that two sets of instructions are provided for getting information about other agents in the world, the "friend" instructions and the "other" instructions. Each "friend" instruction operates on the closest agent having a color

Table 1. Push instructions available for use in SWARMEVOLVE 2.0 agent programs.

| Instruction(s) | Description |
|--|--|
| DUP, POP, SWAP, REP, =, NOOP, PULL, PULLDUP, CONVERT, CAR, CDR, QUOTE, ATOM, NULL, NTH, +, *, /, >, <, NOT, AND, NAND, OR, NOR, DO*, IF | Standard Push instructions (See [14]) |
| VectorX, VectorY, VectorZ, VPlus, VMinus, VTimes, VDivide, VectorLength, Make-Vector | Vector access, construction, and manipulation |
| RandI, RandF, RandV, RandC | Random number, vector, and code generators |
| SetServoSetpoint, SetServoGain, Servo | Servo-based persistent memory |
| Mutate, Crossover | Stochastic list manipulation (parameters from stacks) |
| Spawn | Produce a child with code from code stack |
| ToFood | Vector to energy source |
| FoodIntensity | Energy of energy source |
| MyAge, MyEnergy, MyHue, MyVelocity, MyLocation, MyProgram | Information about self |
| ToFriend, FriendAge, FriendEnergy, FriendHue, FriendVelocity, FriendLocation, FriendProgram | Information about closest agent of similar hue |
| ToOther, OtherAge, OtherEnergy, OtherHue, OtherVelocity, OtherLocation, OtherProgram | Information about closest agent of non-similar hue |
| FeedFriend, FeedOther | Transfer energy to closest agent of indicated category |



Figure 4. A view of SWARMEVOLVE 2.0 in which energy sources shrink as they are consumed and agents are “fatter” when they have more energy.

similar to the acting agent (currently defined as having a hue within 0.1 in a circular hue scale that ranges from 0.0 to 1.0). Each “other” instruction operates on the closest agent having a color that is *not* similar to the acting agent.¹⁴ In some cases, in particular when an agent sets its color once and never changes it, the “friend” instructions will be likely to operate on relatives, since presumably these relatives would set their colors similarly. But since agents can change their colors dynamically each time-step, a “friend” is not necessarily a relative and a relative is not necessarily a “friend.” The term “friend” here should be taken with a grain of salt; the friend/other distinction provides a way for agents to distinguish among each other based on color, but they may use this capability in a variety of ways. Color-based discrimination of other agents, as used here, is similar to the “tag”-based discrimination explored by [3] and [12], although here each agent can change its own tag arbitrarily.

SwarmEvolve 2.0 is an “autoconstructive evolution” system, in which agents are responsible for producing their own offspring and arbitrary reproductive mechanisms may evolve [14]. Whenever an agent attempts to produce a child (by executing the Spawn instruction), the top of its code stack is examined. If the expression is empty (which happens rarely once the system has been running for some time) then a newly generated, random program is used for the child. If the expression is not empty then it is used as the child’s program, after a possible mutation. The probability of mutation is also determined by the parent. A random number is chosen from a uniform distribution from zero to the absolute value of the number on top of the Integer stack; if the chosen number is zero then a mutation is performed. The mutation operation is similar to that used in traditional genetic programming: a random sub-expression is replaced with a newly generated random expression. Note that the program access instructions provide the means for agents to produce their children asexually or sexually, potentially using code from many “mates.”

At the beginning of a SWARMEVOLVE 2.0 run most of the agents, which will have been generated randomly, will not have programs that cause them to seek food and produce offspring; they will therefore die rather quickly and the population will plummet. Whenever the population drops below a user-defined threshold the system injects new random agents into the world. With the parameters used here, however, it usually takes only a few hundred time-steps before “reproductive competence” is achieved—at this point the population is self-sustaining as there are a large number of agents capable of reproducing.

SWARMEVOLVE 2.0 is a complex program with many parameters, not all of which can be addressed in the scope of this short article. However, the source code for the system (including the parameters used in the experiments described below) is available on-line.¹⁵ An animation of a typical action sequence can also be found on-line.¹⁶

6. Emergence of collective behavior in SWARMEVOLVE 2.0

The last two instructions listed in Table 1, FeedFriend and FeedOther, provide a means for agents to transfer energy to one another (to share food). Each of these instructions transfers a small increment of energy (0.01 out of a possible total of 1.0), but only under certain conditions which we varied experimentally (see below). Ordinarily, the use of these instructions would seem to be maladaptive, as they decrease the energy of the acting agents. The use of a Feed instruction thereby makes the feeding agent both more vulnerable and less likely to produce children.

310 Might there nonetheless be some circumstances in which it is adaptive for agents to feed
311 one another? We set out to investigate this question by conducting runs of SWARMEVOLVE
312 2.0 and monitoring the proportion of agents that feed or attempt to feed other agents.¹⁷
313 Because the Feed instructions will occasionally occur in randomly generated code and in
314 mutations, we expect every run to produce some number of calls to these instructions. We
315 expect, however, that the proportion of food sharing agents, when averaged over a large
316 number of runs, will reflect the extent to which food sharing is adaptive.

317 We hypothesized, for example, that dynamic, unstable environments might provide a
318 breeding ground for altruistic feeding behavior. We reasoned as follows, from the perspective
319 of a hypothetical agent in the system: “If the world is stable, and everyone who’s smart
320 enough to find food can reliably get it, then I should get it when I can and keep it to myself. If
321 the world is unstable, however, so that I’ll sometimes miss the food despite my best efforts,
322 then it’d be better for me if everyone shared. Food from others would help to buffer the
323 effects of chance events, and I’d be willing to share food when I have it in order to secure
324 this kind of insurance.” Of course one shouldn’t put too much faith in such hypothetical
325 stories, but they can sometimes be a guide for intuitions. In the present case they led us
326 to conduct a simple experiment in which we varied the stability of the energy sources and
327 the sharing conditions in SWARMEVOLVE 2.0 and measured the proportion of food-sharing
328 agents that resulted.

329 We conducted a total of 1,625 runs under a variety of stability and sharing conditions. We
330 used values of the “stability” parameter ranging from 20 (unstable) to 2,000 (highly stable).
331 The stability parameter governs the frequency with which energy sources begin to drift to
332 new, random locations; the probability that a particular energy source will begin to drift to a
333 new location in any particular time step is $\frac{1}{\text{stability}}$. We collected data on four different sharing
334 conditions. In all of the conditions the potential recipient is the closest agent of similar or
335 dissimilar color, depending on whether the agent is executing the FeedFriend or FeedOther
336 instruction respectively. In all cases the feeding is conditional on the recipient having less
337 energy than the provider. In “waste” sharing the energy is all lost in the transfer, and the
338 recipient receives nothing; we included this sharing condition as a control. In “charity”
339 sharing the recipient receives all of the energy, regardless of whether or not the recipient
340 itself shares energy. In “mutual” sharing the recipient receives all of the energy if it is
341 transferred, but it is transferred only if the recipient has itself executed a sharing instruction
342 at least once in its life. Note that “charity” and “mutual” sharing, unlike the altruistic sharing
343 explored in some other lines of research (e.g., [12]), are *zero sum* transactions: the cost to
344 the donor is equal to the benefit to the recipient. Finally, in “noop” sharing no energy is
345 transferred or lost; this is another control.

346 We collected data only from runs with at least 5,300 consecutive time-steps of repro-
347 ductive competence—there were 936 runs meeting this condition. For qualifying runs we
348 then collected data over the last 5,000 of time-steps, divided into 100-time-step “epochs.”
349 At each epoch boundary we took a census, recording the prevalence of sharing agents
350 normalized by population size. We computed the sum of the number of living agents that
351 had executed FeedFriend at least once and the number of living agents that had executed
352 FeedOther at least once, and then divided this sum by the total number of living agents.
353 Because it is possible for some agents to be counted twice (if they execute both FeedFriend
354 and FeedOther) this is not quite equivalent to the proportion of the population that shares,
355 but it is nonetheless a population-normalized measure of the prevalence of sharing agents.

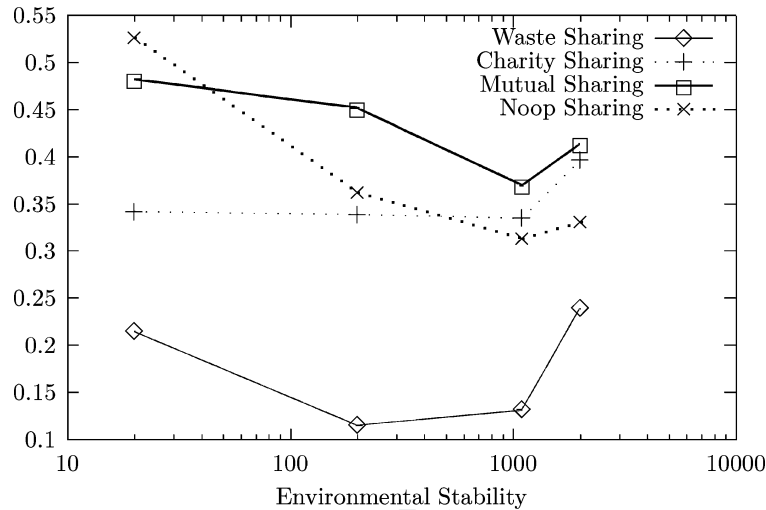


Figure 5. Prevalence of agents that share food (on the y axis) graphed vs. environmental (energy source) stability (on the x axis) for four sharing conditions (see text).

Our results are graphed in Figure 5. There is substantial variance in the data, and the statistical significance of some of the differences visible in the graph is questionable. The distributions are far from Gaussian, so we used the unpaired Wilcoxon two sample test to assess significance. With respect to the vertical differences in the graph, this indicates that all differences at stability 20 are significant ($p \leq 0.015$) except that between “noop” and “mutual” sharing. At stability 200 all differences are significant ($p \leq 0.025$) except for the differences between “noop” and “mutual” and between “noop” and “charity.” At stability 1100 only “waste” is significantly different from the others ($p \leq 0.00047$). At stability 2000 the only clearly significant difference is between “waste” and “mutual” ($p \leq 0.0079$), although the differences between “waste” and the other conditions approach significance ($p \leq 0.09926$ and $p \leq 0.06726$).

The “waste” sharing control produced less food sharing than all of the other sharing conditions; this is what one would expect, as “waste” sharing has costs but no possible benefits. The “noop” sharing control, on the other hand, which has no costs and no benefits, produced more sharing than all other conditions at low stability (although the difference from “mutual” sharing was not significant). The amount of sharing in the “charity” and “mutual” conditions (in which sharing incurs certain costs while providing only potential benefits) was greater than, or at least not significantly less than, the amount of sharing in the “noop” control under several stability settings. This is evidence that altruistic feeding behavior is adaptive in the environment under study. At low stability settings (20 and 200) significantly more sharing occurred in the “mutual” condition than in the “charity” condition, demonstrating that the ability to test for mutuality does indeed have value.

Our hypothesis that dynamic, unstable environments might provide a breeding ground for altruistic feeding behavior was only partially confirmed; while the most food sharing does indeed occur in the least stable environments, the most stable environments also support significant food sharing. To the extent that the hypothesis is confirmed it is interesting to note




Figure 6. A snapshot from a version of SWARMEVOLVE 2.0 in which pairs of agents that have shared energy are connected by black lines, making networks of altruistic energy sharing visible.

382 that a similar effect, involving the preference for cooperation in unpredictable environments,
383 has been observed in a radically different, game-theoretic context [2].

384 These simulations provide a rich framework for investigating the relations between col-
385 lective behavior and evolution, which we have only begun to explore. Some of our more
386 recent explorations have been aided by BREVE's dynamic line-drawing facilities to visualize
387 sharing networks, as shown in Figure 6.

388 7. Conclusions and future work

389 The emergence of collective behavior is an intriguing and at times counter-intuitive phe-
390 nomenon, an understanding of which will have significant impacts on the study of living
391 systems at all levels, from symbiotic microbes to human societies. The work presented
392 in this article demonstrates that new artificial life technologies provide new tools for the
393 synthetic investigation of these phenomena, complementing the well-established analytic
394 methods of evolutionary biology and behavioral ecology.

395 In particular we demonstrated the emergence of a simple form of multicellular organi-
396 zation in evolving populations of agents based on a traditional flocking algorithm. We also
397 demonstrated the emergence of altruistic feeding behavior in a system that is considerably
398 less constrained, as the agents are controlled by evolved computer programs. We believe that
399 this latter system provides significant new avenues of study by allowing agents of arbitrary
400 complexity to evolve within complex, dynamic worlds.

401 Our own plans for this work in the near future include a systematic exploration of the
402 effects of various parameter changes on the emergence of collective behavior. We are making

the source code for SWARMEVOLVE 1.0 and 2.0 freely available in the hopes that others will also contribute to this process; see <http://hampshire.edu/lspector/swarmevolve-1.0.tz> and <http://hampshire.edu/lspector/swarmevolve-2.0.tz>.

Acknowledgments

Raymond Coppinger, Rebecca Neimark, Wallace Feurzeig, Oliver Self-ridge, and several anonymous reviewers provided comments that helped to improve this work. This effort was supported by an NSF Director's Award for Distinguished Teaching Scholars (to Spector), by NSF grant EIA-0216344, and by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30502-00-2-0611.

Notes

1. The use of the term "tag" in this context derives from [3].
2. A system that appears to be similar in some ways, though it is based on 2D cellular automata and the Santa Fe Institute SWARM system, is described at <http://omicrongroup.org/evo/>.
3. Available from <http://www.spiderland.org/breve>
4. <http://www.swarm.org/index.html>
5. <http://repast.sourceforge.net/>
6. The version of [4] in the printed proceedings is actually an early draft, included as a result of a clerical error; please see <http://www.spiderland.org/breve/breve.pdf> for the correct final version.
7. If the resulting velocity would exceed a pre-specified maximum then its magnitude is reduced to the maximum.
8. "Species" here are simply imposed, hard-coded distinctions between groups of agents, implemented by filling "species" slots in the agent data structures with integers ranging from 0 to 2. This bears only superficial resemblance to biological notions of "species."
9. <http://hampshire.edu/lspector/swarmevolve-1.0.tz>
10. The choice to have death and rebirth happen in the same location facilitated, as an unanticipated side effect, the evolution of the form of collective behavior described below. In SWARMEVOLVE 2.0, among many other changes, births occur near parents.
11. Birth energies are typically chosen to be random numbers in the vicinity of half of the maximum.
12. <http://hampshire.edu/lspector/swarmevolve-ex1.mov>
13. <http://hampshire.edu/lspector/swarmevolve-ex2.mov>
14. If there are no other agents meeting the relevant criterion then each of these instructions operates on the acting agent itself.
15. <http://hampshire.edu/lspector/swarmevolve-2.0.tz>
16. <http://hampshire.edu/lspector/swarmevolve2-ex1.mov>
17. For the analyses presented below we did not distinguish between FeedFriend and FeedOther executions; we explored the distinction briefly but there were no obvious patterns in the data.

References

1. J. T. Bonner, *The Evolution of Complexity by Means of Natural Selection*, Princeton University Press: Princeton, NJ, 1988.
 2. A. Eriksson and K. Lindgren, "Cooperation in an unpredictable environment," in *Proc. Eighth Intl. Conf. on Artificial Life*, The MIT Press: Cambridge, MA, 2002, pp. 394–399.
 3. J. H. Holland, *Hidden Order: How Adaptation Builds Complexity*, Perseus Publishing, 1996.
 4. J. Klein, "BREVE: A 3D environment for the simulation of decentralized systems and artificial life," in *Proc. Eighth Intl. Conf. on Artificial Life*, The MIT Press: Cambridge, MA, 2002, pp. 329–334.
- <http://www.spiderland.org/breve/breve-klein-alife2002.pdf>

- 448 5. J. R. Krebs and N. B. Davies, *An Introduction to Behavioural Ecology*, 3rd edition, Blackwell Scientific
449 Publications Ltd.: Oxford, 1981.
- 450 6. J. M. Smith and E. Szathmáry, *The Origins of Life*, Oxford University Press: Oxford, 1999.
- 451 7. H. R. Pulliam and T. Caraci, "Living in groups: Is there an optimal group size?" in *Behavioral Ecology: An*
452 *Evolutionary Approach*, 2nd edition, J. R. Krebs and N. B. Davies (Eds.), Blackwell Scientific Publications
453 Ltd.: Oxford, 1984, pp. 122–147.
- 454 8. T. S. Ray, "Is it alive or is it GA?," in *Proc. Fourth Intl. Conf. on Genetic Algorithms*, Morgan Kaufmann:
455 San Mateo, CA, 1991, pp. 527–534.
- 456 9. M. Resnick, *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*, The MIT
457 Press: Cambridge, MA, 1997.
- 458 10. C. W. Reynolds, "An evolved, vision-based behavioral model of coordinated group motion," in *From Animals*
459 *to Animats 2: Proc. Second Intl. Conf. on Simulation of Adaptive Behavior*, The MIT Press: Cambridge, MA,
460 1993, pp. 384–392.
- 461 11. C. W. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," *Computer Graphics*, vol. 24,
462 no. 4, pp. 25–34, 1987.
- 463 12. R. L. Riolo, M. D. Cohen, and R. Axelrod, "Evolution of cooperation without reciprocity," *Nature*, vol. 414,
464 pp. 441–443, 2001.
- 465 13. R. M. Silby, "Optimal group size is unstable," *Anim. Behav.*, vol. 31, pp. 947–948, 1983.
- 466 14. L. Spector and A. Robinson, "Genetic programming and autoconstructive evolution with the push program-
467 ming language," *Genetic Programming and Evolvable Machines*, vol. 3, no. 1, pp. 7–40, 2002.
- 468 15. L. Spector, "Adaptive populations of endogenously diversifying pushpop organisms are reliably diverse," in
469 *Proc. Eighth Intl. Conf. on Artificial Life*, The MIT Press: Cambridge, MA, 2002, pp. 142–145.
- 470 16. N. Zaera, D. Cliff, and J. Bruten, "(Not) Evolving collective behaviours in synthetic fish," in *From Animals to*
471 *Animats 4: Proc. Second Intl. Conf. on Simulation of Adaptive Behavior*, The MIT Press: Cambridge, MA,
472 1996, pp. 635–644.