Exploratory Design Of Swarms

Eric Bonabeau, Pablo Funes, Belinda Orme
Icosystem Corporation, 10 Fawcett Street, Cambridge, MA 02138, USA.
Corresponding author: eric@icosystem.com

Abstract

In order to fulfill the true promise of swarm intelligence, and more generally of decentralized, self-organizing intelligence, a major design problem has to be overcome. Designing the individual-level rules of behavior and interaction that will produce a desired collective pattern in a group of human or non-human agents is difficult because the group's aggregate-level behavior may not be easy to predict or infer from the individuals' rules. While the forward mapping from micro-rules to macro-behavior in self-organizing systems can be reconstructed using computational modeling techniques, the inverse problem of finding micro-rules that produce interesting macro-behavior poses significant challenges, all the more as what constitutes "interesting" macro-behavior may not be known ahead of time. An exploratory design method is described in this paper. It relies on interactive evolution. We show how it can be used to discover new, "interesting" patterns of collective behavior when one does not know in advance what the system is capable of doing, a generic situation in swarm design.

Keywords: collective behavior, swarms, design, interactive evolution

1 Introduction

Designing the individual-level rules of behavior and interaction that will produce a desired collective pattern in a group of human or non-human agents is difficult because the group's aggregate-level behavior may not be easy to predict or infer from the individuals' rules. For example, the aggregate-level properties of a traffic jam (Helbing *et al.*, 2000), a crowd evacuating a public space (Still, 1993) or the stock market (Palmer *et al.*, 1994) cannot easily be derived from knowing the rules of behavior and interaction of drivers, people or investors. Agent-based modeling (ABM) (Reynolds, 1987; Epstein & Axtell, 1996; Axelrod, 1997; Bonabeau, 2000, 2002), or micro-simulation, is often the only way to capture the emergent properties resulting from the behavior and interactions of the group's constituent units or agents, particularly but not only when the individual-level rules are discrete. While ABM is useful in producing aggregate-level patterns from individual-level rules, finding the appropriate rules still requires manual search and tinkering when (1) the collective-level patterns may be difficult to formalize into a mathematical detector and therefore the evaluation of a solution cannot be automated, and/or (2) the collective-level patterns made possible by the individual-level rules are not even known ahead of time. If the desired collective-level pattern is known and its detection can be automated, then traditional search and optimization algorithms can be utilized. This paper focuses on cases where (1) and/or (2) is true. We show that in such cases a technique originally developed to generate "interesting" images and pieces of art (Dawkins, 1987; Sims, 1991, 1992, 1993) can be used to design the individual-level rules of behavior and interaction to produce "interesting" collective-level patterns.

The technique (see Takagi, 2001 for a review) is a directed search evolutionary algorithm which requires human input to evaluate the fitness of a collective-level pattern (here, the fitness might be how close the collective-level pattern is to the desired pattern, or how interesting the pattern is) and uses common evolutionary operators such as mutation and crossover (Forrest, 1993) to breed the individual-level rules that produced the fittest collective-level patterns. Using a simple example of a human game that can be played in small groups (Bonabeau, 2002; Bonabeau *et al.*, 2003; Funes *et al.*, 2003), we show that this approach is particularly powerful as an exploratory design technique, when the aggregate-level capabilities of the

system are not known. Interactive evolutionary computation (IEC), as this technique is known, combines computational search with human evaluation (Takagi, 2001).

Some of the individual-level rules discovered using IEC are presented together with their corresponding striking, unexpected patterns. In itself, the discovery of the rules is important as it shows that it is possible to design simple rules to produce robust (genetic robustness being a by-product of the evolutionary method: to be discovered, a fitness peak has to be reasonably stable under a number mutations), collective-level patterns; in addition, the IEC technique used to discover these rules is very generic and makes it possible to systematically discover novel phenomena in self-organizing systems.

2 The aggressor-defender game

To illustrate the approach, a game of aggressors and defenders is used. Two rules can be used by the players:

- Rule #1: Pick two people A and B in the group. Then start moving in such a way as to always put B (your defender) between yourself and A (your aggressor).
- Rule #2: Pick two people A and B in the group. Then start moving in such a way as to always put yourself between A (the aggressor) and B (the defendee).

If for example everyone in the group follows Rule #1 and picks A and B randomly, the resulting collective-level pattern is chaotic, sometimes room-filling motion across the room, often constrained by the walls, which leads to some wall following. If on the other hand everyone in the group follows Rule #2 and picks A and B randomly, the whole group collapses onto a single cluster. These two simple versions of the game can easily be played with a group of 8 or more people (our experience includes playing the game with up to 400 people) and produce aggregate-level patterns that cannot easily be predicted from an examination of the individual rules. For example, in both versions there is exactly the same number of aggressors, defenders, defendees. While it is difficult to predict the outcome of even the simplest versions of the game, Anderson (2003) [p. 9, this volume] has established elements of a mathematical proof for the two situations described above (all participants following the same rule and picking their A's and B's randomly). A complementary approach consists of simulating the mapping from micro-rules to macro-behavior through agent-based modeling (ABM) (Epstein & Axtell, 1996; Bonabeau, 2002). Figure 1 illustrates the patterns observed in the two simple cases described above when simulated using agent-based modeling. See Funes et al. (2003) for more detail on the ABM.

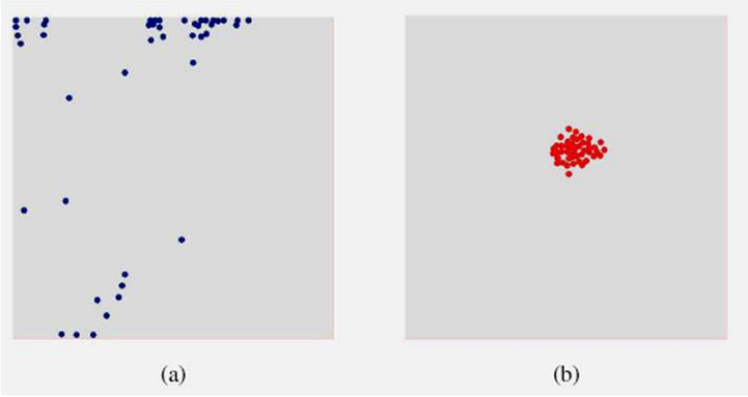


Figure 1: Micro-rules #1 (a) and #2 (b) lead to two dramatically different collective behavior in reality and in the predictive agent-based simulation.

In some situations other than the two simple ones described above, we discovered by playing the game live that the aggregate-level pattern could be quite different from the above-mentioned patterns. For example, if by chance everyone picks the same person as aggressor or defender of defendee, the resulting aggregate-level pattern is very different. In other situations the relationship graph has several disconnected components (although the probability of observing more than one component if participants pick A and B uniformly randomly tends to 0 as n^{-3}; C. Anderson, unpubl. ms), leading to the formation of several clusters. After these phenomena were observed in vivo, the corresponding micro-rules were simulated using ABM and could be re-created *in silico*, proving the predictive power of ABM. The observation of these unexpected collective patterns of behavior prompted us to ask the following question: would it be possible to design social networks (characterized by relationship graphs: who interacts with whom and what is the nature of the interaction?), instead of creating random ones, to produce interesting aggregate-level patterns, with no *a priori* knowledge of what interesting patterns this system could produce?

More precisely, every individual is characterized by the following set of rules or properties:

- Does the individual follow Rule #1 or Rule #2?
- Who is A?
- Who is B?

The size of the space of relationship graphs (which we will also call rule space) grows fast with the number of individuals N. While rule space is large, it is likely that most resulting collective patterns will either be random-looking or reducible to one of the two basic patterns described above (chaotic behavior or clustering). However we do not know what such a system can or cannot do, that is, we do not know what kind of collective-level patterns to expect other than chaotic motion and clustering. In Section 3 we describe a method to search for the relationship graphs, or individual-level rules, that will produce the most interesting collective-level patterns with only a vague notion of what interesting means: neither random nor reducible to one of the two basic patterns.

3 Interactive evolution

3.1 Search mechanism

The IEC search method works as follows. A small initial population of relationship graphs is generated. The resulting collective-level dynamical patterns are generated using an agent-based simulation and presented to a human observer. The observer selects the collective-level patterns that are the most interesting—the fittest individuals in the population according to whatever set of objective and subjective criteria the observer may be using. A new population (new generation) of relationship graphs is generated by applying mutation and crossover operators to the relationship graphs that correspond to the previous generation's fittest collective-level patterns (Forrest, 1993). The new population is then simulated and the resulting collective-level patterns presented to the observer, and so forth. This procedure is iterated until interesting patterns emerge from the search.

The user interface, which is a critical component of the method as it is based on visualizing the solution that the observer evaluates solutions, is shown in figure 2. Obviously this method can only work if the population size is kept small and if interesting patterns emerge after a reasonably small number of generations. Whether or not the method can work depends on the nature of the fitness landscape and on the dimensionality of the problem. For example, a highly rugged fitness landscape or a flat landscape is unlikely to lend itself to this type of search method (Theraulaz & Bonabeau, 1995; Bonabeau *et al.*, 2000).

3.2 Details of the evolutionary algorithm

The evolutionary algorithm used at every generation replaces the current generation of games be either recombination or mutation of the selected game. The notation we use to pertain to a given generation is

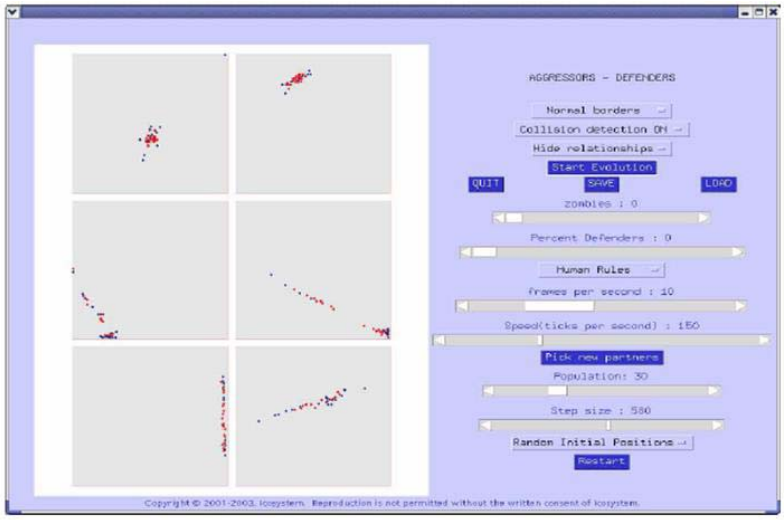


Figure 2: User interface with six playgrounds. Clicking on the playground that displays the user's preferred behavior results in the remaining five becoming mutants or recombinations of the selected one, as described in section 3.2.

the subscript notation such that game i in generation n and $n+1$ is given by i_n and i_{n+1} respectively.

Selection of the fittest game

One game, j_n, is selected to be preserved in generation $n+1$ from the previous generation, n. The remaining games, i_{n+1} $(i \neq j)$, will consist of agents who are derived from either a recombination or a mutation of some the agents in the original game, j_n. This implies mutation and recombination occur with a probability of 0.5.
(For example: If we are evolving 9 games, on average: 1 will be identical to the game chosen in the previous generation; 4 will have some agents in their population altered by recombination; and 4 will have some agents in their population altered by mutation.)

Recombination

If the rules for game i_{n+1} are to be created from a recombination of the rules describing agents in game j_n, then up to 10% of the total agents in game i_{n+1} are derived in the following manner. A typical agent selected to be changed is called the target agent p. A source agent, q, is chosen randomly from the population to provide the rules to replace some of the rules in the target agent p. The rules which can be substituted from q into p are any of

1. Person A,
2. Person B,
3. Rule of the game to be played (# 1 or #2).

Each person or agent in the game has their own set of these three rules. In recombination, while the source agent's rules remain unchanged some of the target agent's rules can be replaced by those from the source agent. The number of rules from p which are substituted for those in q is chosen randomly.

Mutation

If the rules for game i_{n+1} are to be created from a mutation of the agents' rules in game j_n, then up to 10% of the total number of agents in game i_{n+1} are derived in the following manner. To mutate an agent's rules, three possible mutations can occur based on the three rules for each of the agents. The evolution randomly chooses, with equal probability, to mutate either

1. Person A,
2. Person B,
3. Rule of the game to be played (# 1 or # 2).

Each mutation has a probability of one third of occurring.

4 Results

A number of patterns discovered using the IEC technique described above are presented in this section. Figure 3 shows several such patterns.

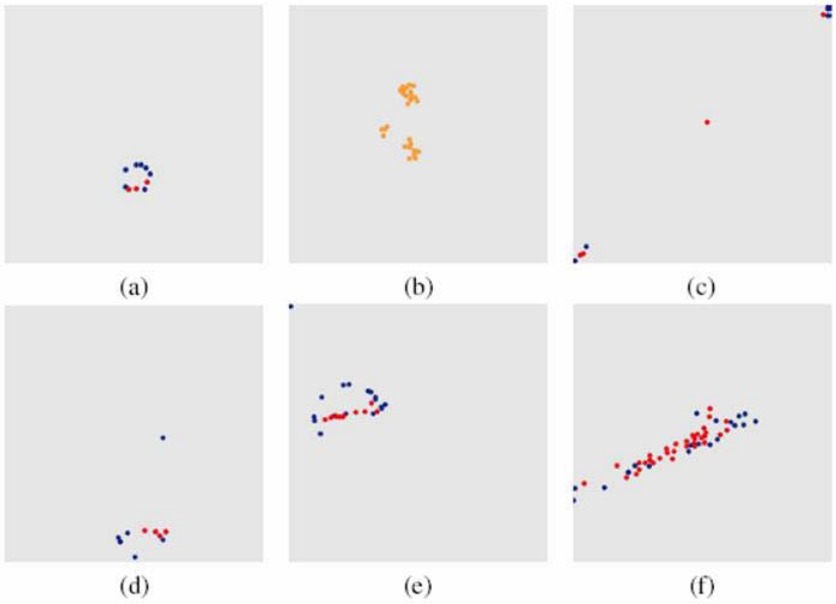


Figure 3: A few examples of evolved behaviors. (a) Circle: agents chase each other around in a circle. (b) Juggle: two blobs fuse and reemerge and sometimes toss a smaller blob at each other. (c) Corner-middle: two groups of agents go to opposite corners while one stays in the middle. (d) Pursuer-evaders: an agent follows a larger group that slows down, is reached by the pursuer, then escapes again. (e) Chinese streamer: a D shape that moves around. (f) Somersault: a thick line that makes a 360 degree turn, then stops, then turns back in the opposite direction.

Figure 4 shows snapshots of the spatio-temporal dynamics of one of the patterns, the Chinese streamer (CS). The CS pattern is particularly interesting in that it is unexpected (it was impossible to predict that the system could display this type of behavior under the right relationship graph), robust to mutations

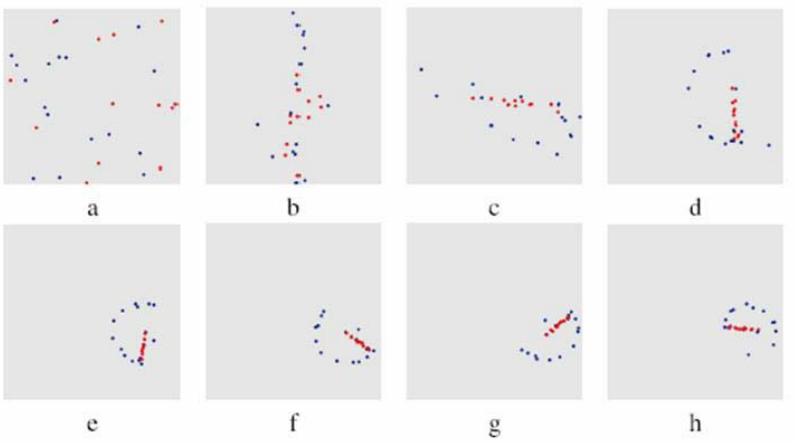


Figure 4: "Chinese streamer" pattern. From a random initial placement, a pattern quickly emerges (a-d) and starts turning, stabilizing in a shape with a handle and trailing ribbon which rotates smoothly. The direction of rotation can be clockwise or counterclockwise (as here), presumably depending on the initial positions.

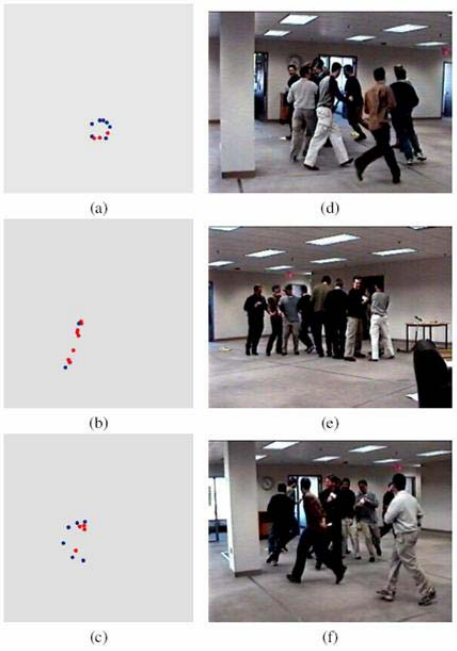


Figure 5: Three rules designed for swarms of 10 agents, evolved using the IEC interface (a-c) and then given to a group of people (d-f). Rule "circle" (a,d makes all agents run around in a circle; rule "align" (b,e) made them form a straight line and rule "Chinese streamer" resulted in a central cluster with a tail or "ribbon" circling behind.

(a by-product of evolutionary search) and totally insensitive to initial conditions (when the appropriate relationship graph is in place, the Chinese streamer always forms regardless of the initial positions of the participants; initial conditions influence the rotation of the pattern, which can be clockwise or counter-clockwise depending on the details of the participants' locations). Lastly, Figure 5 illustrates the ultimate real-world test of the approach: when the rules discovered by applying IEC to the dynamic output of the simulation are given to humans, the predicted patterns do emerge.

5 Conclusions

We have illustrated with a simple game example how to tackle the swarm design issue. Using this approach requires a shift in mindset from the traditional top-down design approach: here, design is exploratory because the aggregate-level capabilities of the system are not known ahead of time. Some of the patterns discovered in the simple game system were extremely surprising and could not be foreseen by any human engineer. Obviously the patterns discovered may sometimes be difficult to understand and need to reverse-engineered. For example the CS pattern has been reverse-engineered and is now well understood. This approach to designing self-organizing systems has a wide range of applications, from collective robotics to distributed control. One example: radio-frequency tags, known as RFIDs. Although RFIDs have recently become very popular, most users intend to use them with a centralized mindset without knowing what a swarm of RFIDs might be able to do collectively. Our exploratory design approach enables an open-minded search for the hidden capabilities of such a system.

References

Anderson, C. 2003. Linking micro- to macro-level behavior in the aggressor-defender-stalker game. Pages 9–16 in: *Proceedings of the Second International Workshop on the Mathematics and Algorithms of Social Insects* (C. Anderson & T. Balch, eds.), Georgia Institute of Technology.

Axelrod, R. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration.* Princeton University Press, Princeton, NJ.

Bonabeau E. 2000. Business applications of agent-based simulation. *Adv. Complex Syst.* **3**: 451–461

Bonabeau, E. 2002. Agent-based modeling: methods and techniques for simulating human systems. *Proc. Nat. Acad. Sci. USA* **99**: 7280–7287

Bonabeau, E., Guérin, S., Snyers, D., Kuntz, P., Theraulaz, G. & Cogne, F. 2000. Complex three-dimensional architectures grown by simple agents: an exploration with a genetic algorithm. *BioSystems* **56**: 13–32.

Bonabeau, E., Hunt, C.W., & Gaudiano P. 2003. Agent-based modeling and designing novel decentralized command and control systems paradigms. Presented under the "Modeling and Simulation and Network- Centric Application Topics" for the *8th International Command and Control Research and Technology Symposium*, June 17-19, 2003, National Defense University, Washington, DC.

Dawkins. R. 1987. *The Blind Watchmaker.* W. W. Norton, New York.

Epstein J. M., Axtell R. L. 1996. *Growing artificial societies: social science from the bottom up.* MIT Press, Cambridge, MA.

Forrest, S. 1993. Genetic algorithms: Principles of adaptation applied to computation. *Science* **261**: 872–878.

Funes, P., Orme, B. & Bonabeau, E. 2003. Evolving emergent group behaviors for simpe humans agents. Pages 76-89 in: *7th European Conference on Articial Life (ECAL 2003): Workshop and Tutorials* (P. Dittrich, J.T. Kim, eds.), Dortmund, 14-17 September, 2003.

Helbing, D., Farkas, I. & Vicsek, T. 2000. Simulating dynamical features of escape panic. *Nature* **407**: 487–490.

Palmer, R. G., Arthur, W. B., Holland, J. H., Le Baron, B. & Tayler, P. 1994. Artificial economic life: a simple model of a stockmarket, *Physica D* **75**: 264–274.

Reynolds, C. 1987. Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics* **21**: 25–34.

Sims, K. 1991. Artificial evolution for computer graphics. *Computer Graphics* **25**: 319–328.

Sims, K. 1992. Interactive evolution of dynamical systems. Pages 171-178 in: *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life* (F. J. Varela & P. Bourgine, eds.), MIT Press, Cambridge, MA.

Sims, K. 1993. Interactive evolution of equations for procedural models. *Vis. Comput.* **9**: 446–476.

Still, K. G. 1993. New computer system can predict human behaviour response to building fires. *Fire* **84**: 40–41.

Takagi, H. 2001. Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proc. IEEE* **89**: 1275–1296.

Theraulaz, G. & Bonabeau, E. 1995. Coordination in distributed building. *Science* **269**: 686–689.