A Knowledge-Tracking Algorithm For Generating Collective Behavior In Individual-Based Populations

Robert W. Schmieder

Sandia National Laboratories
Livermore, CA 94551 USA
rwschmi@ca.sandia.gov

Abstract

An algorithm is described that tracks the amount and distribution of mutual knowledge within a population of interacting individuals. The dynamics of the individuals are modified by their knowledge of other individuals, thus producing emergent collective behavior of the population. The algorithm was designed for populations of minimally knowledgeable objects such as simple animals, but it also provides a conceptual framework for describing the dynamics of populations of inanimate objects. The use of the amount and distribution of knowledge as a dynamical principle, and its relationship to information theory and other dynamical theories, are presented.

Introduction

This work was motivated by recent progress on individual-based ecosystems (DeAngelis, 1992) and collective behavior of populations of simple objects (Langton, 1987). The central challenge in such simulations is to carry a system of interacting individuals forward in time, in order to assess the effects of parameter variations. These problems are sometimes described as information-driven: the size of a tractable problem and the speed of the algorithm are limited by the enormous amounts of information that must be processed.

The approach adopted here is to focus on *knowledge* among the individuals of the population. We have a general concept of "knowing," and could say that each individual in a population in some sense "knows" about other individuals. Clearly, individual dynamics will depend strongly on what (or who) that individual knows.

In this work, we present an algorithm (abbreviated KtA) that tracks the amount and distribution of knowledge among identical individuals comprising an isolated population, and a dynamical principle for generating the emergent collective behavior of the population dependent on that knowledge. We do not track the knowledge itself, only a quantity giving the amount and distribution of that knowledge. This quantity is precisely defined in terms of probabilities, hence provides a link to classical and quantum dynamics, statistical mechanics, information theory, and knowledge representation. The KtA approximates the basic dynamics of knowledge in real systems, including sharing and various loss mechanisms.

Knowledge in a Population of Individuals

We begin by assuming that the individuals in the population are capable of carrying along with them a quantity of "knowledge" about other individuals. We do not require individuals to be sufficiently talented to access and process their knowledge in any manner resembling intelligence; the individuals may be nothing more than simple devices capable of only a small number of reflexive actions.

We assume that the individuals interact in pairs according to some specified physical dynamics. If these interactions are infrequent and local, they resemble collisions, and the model is connectionist (Farmer, 1990). If the interactions are continuous and nonlocal, they resemble long-range forces. When an interaction occurs, two things happen: (1) the interacting partners mutually share a fraction of the knowledge they carry; (2) all other individuals lose a fraction of the knowledge their knowledge of the interacting partners.

Knowledge sharing due to interaction is intuitively obvious: Assume that individuals {A,B,C} initially have no knowledge of each other; their behavior is completely independent of each other. After an {A,B} interaction, individuals {A} and {B} know about each other, but neither knows anything yet about {C}. After a {B,C} interaction, individuals {B} and {C} know about each other, and therefore, {C} knows something about {A}, although {A} knows nothing yet about {C}. This process leads to an expansion and general increase of knowledge. While the process is similar to diffusion, it has no direct analog in statistical mechanics.

Knowledge loss due to interaction is *not* intuitively obvious. The following argument demonstrates the origin of this decrease: In any interaction, both partners are inevitably and permanently altered (this must be so, or else no interaction took place), indicated by {A,B}—>{A',B'}. Next, {B'} and {C} interact: {B',C}—>{B",C'}. However, since {A'} has no means to know how {B'} was altered by {C} during this event, it knows less about {B"} than it did about {B'}. With the passage of more and more time, {A} has less and less knowledge of {B} because {A} has no knowledge of {B}'s other interactions. We are familiar with this process in a social context; with the passage of time, we tend to know less and less about a lost friend. This process clearly leads to a decrease in knowledge, and while it bears some similarity to spontaneous decay and dissipative processes, there is again no direct analog in statistical mechanics.

There are, in addition, other causes of knowledge alteration: an individual may simply forget, or the knowledge may get corrupted with other knowledge. In the KtA algorithm to be described in the next section, we will include only the two mechanisms described above: gain through sharing and externally induced loss.

The dynamics of the individuals clearly depend on their mutual knowledge. If the mutual knowledge of the individuals in the population is zero, they all act independently and there is no collective dynamics. If the mutual knowledge is high—most individuals know a great deal about most other individuals—the population is strongly interacting, and exhibits collective behavior. In the limit that the mutual knowledge is total (everyone knows everything about everyone else), the population is fully connected; it behaves like a rigid body. In this case an effect on one individual is fully felt by all members of the population.

The Knowledge-Tracking Algorithm (KtA)

The foregoing ideas are implemented in a knowledge-tracking algorithm (KtA) as follows: For a population of N individuals, we define an NxN matrix K; the matrix element K_{ij} numerically represents how much individual {i} "knows" about some property of individual {j}. The K_{ij} have the constraints listed in Table 1.

Table 1—Constraints on the KtA matrix elements

Self knowledge	$K_{ii}=1$	Every individual always knows everything about itself.
Mutual knowledge	$0 \leq K_{ij} \leq 1$	Individuals know between nothing and everything (inclusive) about all other individuals.

Each row in the matrix represents what one individual knows about all other individuals (at one generation). Each column represents what all other individuals know about one individual. The principal diagonal represents what each individual knows about itself (always =1).

The matrix K evolves as follows: In an interaction {i,j}, the matrix element at generation n is converted according to a rule into the corresponding matrix element for generation n+1: $K_{ij}(n) \longrightarrow K_{ij}(n+1)$. The rule has the form listed in Table 2.

Table 2—Form of the KtA transformation rules

Interacting individuals	$K_{iq}(n+1) = K_{iq}(n) + f[K_{jq}(n)]$ $K_{jq}(n+1) = K_{jq}(n) + f[K_{iq}(n)]$ $f(x)$ = increasing function of x $q=1 \ldots N$	Interacting individuals mutually acquire some of their partners' knowledge
All other individuals	$K_{qi}(n+1) = g[K_{qi}(n)]$ $K_{qj}(n+1) = g[K_{qj}(n)]$ $g(x)$ = decreasing function of x $q=1 \ldots (\neq i,j) \ldots N$	All *other* individuals lose some of their knowledge of the interacting partners

The mechanics of the KtA matrix transformation is visualized as follows: An {i,j} interaction causes rows i and j to be positively mixed (vector addition). Columns i and j (except the matrix elements involving the interacting partners) are decreased equally. Any matrix element that increases to $K_{ij}(n)>1$ is automatically truncated to $K_{ij}(n)=1$, and any element that may fall below $K_{ij}(n)<0$ is held at $K_{ij}(n)=0$. The matrix diagonal is always unity, $K_{ii}(n)=1$.

Several lumped descriptors of the matrix K are useful. The sum of elements in a row represents the total knowledge held by individual {i}:

$$K_i(n) = \Sigma_j K_{ij}(n)$$

The sum of all matrix elements is the total knowledge in the population:

$$K(n) = \Sigma_i K_i(n) = \Sigma_i \Sigma_j K_{ij}(n).$$

The Linear Symmetric Algorithm

Within the constraints given in the previous section, a wide variety of transformation rules is possible. The simplest, and probably the most reasonable physically, is to assume that during interaction each partner acquires a fraction (a) of the knowledge of the other partner, and all individuals other than the interacting partners lose a fraction (b) of their knowledge of both partners. The matrix transformation effected by this algorithm is shown explicitly:

$$\begin{array}{lllllll}
K_{11} & \cdots & (1-b)\,K_{1i} & \cdots & (1-b)\,K_{1j} & \cdots & K_{1N} \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
K_{i1}+aK_{j1} & \cdots & K_{ii} & \cdots & K_{ij}+aK_{jj} & \cdots & K_{iN}+aK_{jN} \\
K_{j1}+aK_{i1} & \cdots & K_{ji}+aK_{ii} & \cdots & K_{jj} & \cdots & K_{jN}+aK_{iN} \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
K_{N1} & \cdots & (1-b)\,K_{Ni} & \cdots & (1-b)\,K_{Nj} & \cdots & K_{NN}
\end{array}$$

where it should be remembered that $K_{ii}=1$ always and that all matrix elements are truncated to 1. This rule is linear and symmetric between collision partners. The algorithm is effected by performing this matrix transformation repeatedly. For example, consider four individuals {**A,B,C,D**} starting with only self-knowledge ($K_{ij}=\delta_{ij}$). If they interact in the sequence (**B**,**C**) (**A**,**B**) (**C**,**D**) (**B**,**D**), the matrix evolves as shown here. Note that K_{ij} are polynomials of lower degree than the generation n.

Generation n—Intereracting pairs

——1——(B,C)—2——(A,B)——3———(C,D)———4———(B,D)————5————

	1	2	3	4	5
A =	{1 ,0 ,0 ,0}	{1 ,0 ,0 ,0}	{1, a,a^2,0}	{1, a,$(1-b)a^2$,0}	{1, $(1-b)a$, $(1-b)^2a^2$, 0 }
B =	{0 ,1 ,0 ,0}	{0 ,1 ,a ,0}	{a, 1, a ,0}	{a, 1,$(1-b)a$,0}	{ a, 1 , $(1-b)a+a,a^2$}
C =	{0 ,0 ,1 ,0}	{0 ,a ,1 ,0}	{0, $(1-b)a$, 1 ,0}	{0, $(1-b)a$, 1 ,a}	{a^2, $(1-b)a+a$, 1 , a }
D =	{0 ,0 ,0 ,1}	{0 ,0 ,0 ,1}	{0, 0, 0 ,1}	{0, $(1-b)a^2$, a ,1}	{0 , $(1-b)^2a^2$, $(1-b)a$, 1 }.

Properties of the Algorithm

The matrix K is not necessarily symmetric. If K *is* symmetric, each individual knows as much about another individual as the other knows about the first. But for large values of a or large number of iterations, significant asymmetries are generated, even starting from a symmetric initial matrix. Physically, this means that individual {i} may know more (or less) about individual {j} than {j} knows about {i}.

The diffusive nature of the KtA is illustrated in Fig. 1, in which we have plotted the non-zero matrix elements $K_{ij}(n)$ as white squares. Initially, a population of 30 individuals with no knowledge of each other was established. The only non-zero matrix elements were the diagonal $K_{ii}(0)=1$. At successive generations, the linear, symmetric algorithm with a=0.5, b=0.5 was applied to the population, allowing only nearest neighbors to share knowledge. After 30 generations, the diagonal has "diffused" into a broad band; each individual knows about several others within its immediate neighborhood, but knows nothing of more distant individuals. After 300 generations the matrix has equilibrated, and never diffuses beyond the ragged diagonal band. This example also illustrates the fact that knowledge "islands" can occur, and exceptional individuals can be produced.

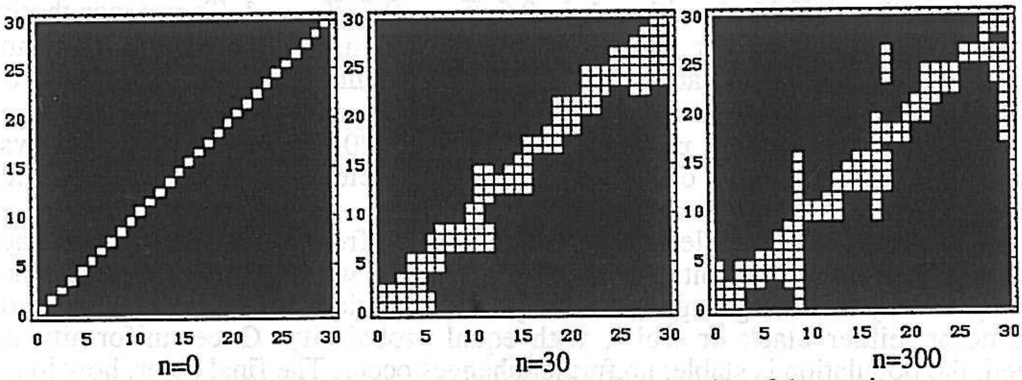

Figure 1—Three frames from a sequence showing the evolution of the matrix elements K_{ij} from initialization to equilibration.

At large times, populations approach an equilibrium total knowledge K_e. Any initial matrix K is gradually transformed into a matrix with $K(n \longrightarrow large) \approx K_e$. The KtA is therefore a *learning* algorithm. Figure 2 shows K(n) for a population of 20 individuals that evolved according to the linear symmetric algorithm with a=1, b=0.1. Initially the individuals were independent, and the total knowledge was 20. As the individuals interacted, their mutual knowledge increased, then reached an equilibrium value K_e slightly below the maximum possible (400). Decreasing the value of a causes the learning transient to be slower. Increasing the value of b causes the equilibrium value K_e to be lower. At equilibrium, the total knowledge K(n) fluctuates chaotically around its equilibrium value K_e. The magnitude of these fluctuations is smaller in a population with a larger number N of individuals.

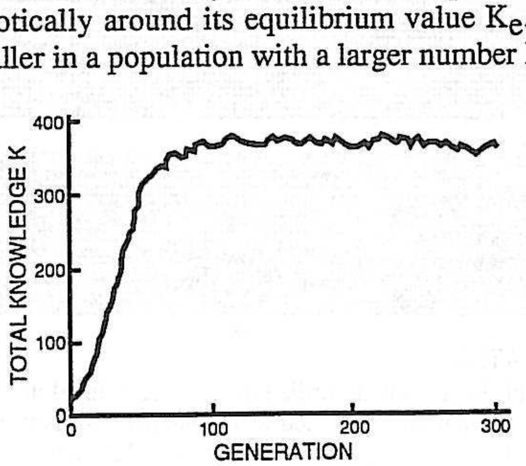

Figure 2 —Growth of the total knowledge K following initialization of a population of 20 individuals with zero mutual knowledge.

Example 1: Imitation

This example implements a population of black and white individuals in a 1-dimensional space. The dynamics allows individuals to interact only with their nearest neighbors and to adopt the color of the neighbor if the individual has a specified knowledge of that neighbor. This algorithm was implemented per Table 3.

Table 3—Rules for the Imitation KtA

Initial configurations are specified arbitrarily.

Interactions are implemented by selecting one partner {i} at random. The other partner is the nearest neighbor, either {i-1} or {i+1}, with equal probability.

When an interaction occurs, the matrix K is updated according to the linear symmetric KtA given above, using fixed values of a and b. Then each partner searches all other individuals to identify the individual about which it knows the most (maximum K_{iq}, K_{jq}). If more than one individual is thereby identified, one is selected randomly from among those. Then, if $K_{min}<K_{iq},K_{jq}<K_{max}$ each partner assumes the color of {q}; otherwise they retain their previous colors.

Figure 3 presents typical results obtained for a population of 20 individuals evolving according to Table 3, with a=1, b=0.5, K_{min}=0.5, K_{max}=1. To produce these data, the K matrix was first evolved for sufficient time to have achieved an equilibrium configuration (it reached K_e≈0.7). Then the initial population was colored at random, and allowed to evolve for 200 (or more) generations. The evolution of the population was represented by plotting the color of the 20 individuals (vertically) vs generation. From a large set of results, we have selected three typical panels to include in Fig. 3.

The following typical collective behaviors emerge from these simulations: The population tends to fragment into groups of similar color which persist for some time but then give way to other groups. Eventually the population is completely converted to one color, either black or white, with equal probability. Once uniformity is achieved, the population is stable: no further changes occur. The final color, how long it takes to achieve stability, and any details about the transient groups, are completely unpredictable from the initial configuration. We emphasize that the dynamics of this KtA are completely deterministic, except that the sequence of collision partners is selected randomly.

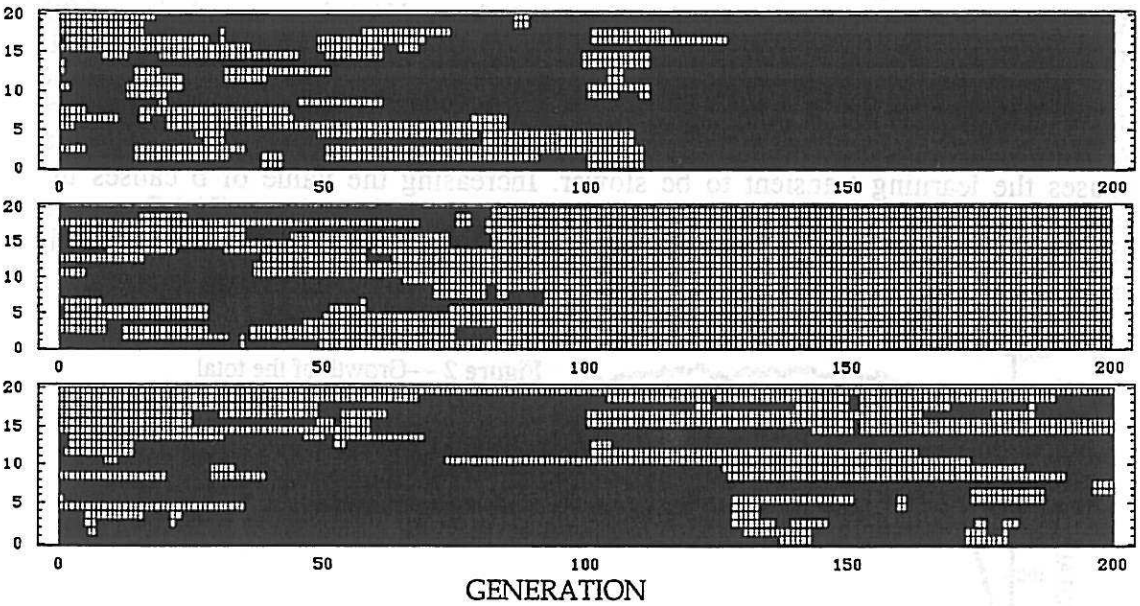

Figure 3—Evolution of a population of 20 individuals interacting according to the rules defined in Table 3. For each of the three panels, the population is restarted with random configuration. The final configuration is always uniformly one color (panel 3 was completed at later times), but that color is unpredictable from the initial configuration.

Example 2: Pursuit-and-Flight

This example implements a population of male and/or female individuals moving within a 2-dimensional box. The dynamics allows some individuals to be attracted to or repelled by other individuals. The model encompasses pursuit and capture, avoidance, clustering, flocking, swarming, stalemating, and a variety of other collective phenomena commonly seen in nature. In addition, upon interaction, the sex of the individuals may be altered, providing the possibility of spatial/sexual coupling.

This example was implemented as described in Table 4.

Table 4—Rules for the Pursuit-And-Flight KtA

Initial configurations are specified arbitrarily.

The positions of individuals vary in time according to

$$x_i(t+\Delta t) = x_i(t) + \Sigma_j K_{ij} (S_{ij} / D_{ij}) X_{ij}(t) \Delta x$$
$$y_i(t+\Delta t) = y_i(t) + \Sigma_j K_{ij} (S_{ij}/D_{ij}) Y_{ij}(t) \Delta y$$

The direction cosines are

$$X_{ij} = (x_i - x_j)/D_{ij}$$
$$Y_{ij} = (y_i - y_j)/D_{ij}$$

and the distances between individuals are

$$D_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \ .$$

and $\Delta x, \Delta y$ are constants.

The state of each individual contains two quantities: sex and position.

K_{ij} represents the individuals' knowledge of sex of other individuals. Sex-changing interactions occur only when one individual approaches another within the sum of their sizes (a "collision"):

$$|x_i - x_j| \leq 2L \ \ \textit{and} \ \ |y_i - y_j| \leq 2L \ ,$$

where L is arbitrarily specified. When this occurs, the matrix K is updated according to the linear symmetric KtA with parameters a,b, and the sexes of partners {i}, {j} are changed with probability b.

S_{ij} accounts for the attraction and/or repulsion between sexes. It is selected to be one of the following four arbitrary values $\{S_{mm}, S_{mf}, S_{fm}, S_{ff}\}$.

The factor (S_{ij} / D_{ij}) represents a pseudoforce that produces motion.

Individuals approaching the boundary of the box are either repositioned a small distance within the box, or reflected from the wall.

The direction cosines serve to orient the attraction or repulsion toward the interaction partners. The factor S_{ij}/D_{ij} in the equations of motion provides a range on the interactions. This factor is quite arbitrary, and was selected mainly to produce interesting motions. It results in strong clustering of individuals and the ability of an individual to be isolated at some distance from the group. We could also have phrased this term as a distance-dependent knowledge K'_{ij} updated at each time step.

The four values $\{S_{mm}, S_{mf}, S_{fm}, S_{ff}\}$ set the qualitative nature of the collective motion. A large negative value for one of these parameters means that individuals of the second sex will be strongly attracted to individuals of the first sex. A large positive value will produce strong corresponding repulsion. A value of zero means that there will be no influence on the motion. A set found to be particularly interesting is {-0.1,1,-1,-0.1}, which produces motion in which males chase females, females flee males, and there is a weak male-male and female-female attraction. We emphasize that the specific dynamics of this KtA are arbitrary; the purpose is to illustrate the effects of mutual knowledge in driving the population behavior. In fact, this form of the KtA produces *no* motion, except by virtue of mutual knowledge.

The basic behavior of populations evolving according the Table 4 is shown in Fig. 4. The population consists of 15 individuals, about half males (circle with dot) and the rest females (open circle). Sex changes are not permitted (b=0). The parameters S={-0.1,1,-1,-0.1} produce the pursuit/flight behavior. In Fig. 4(A), with the mutual knowledge low, the males wander about independently, chasing whatever female happens to be close. In Fig. 4(B), when the mutual knowledge is high, a group of males starts near the center and pursues a group of females in the SE corner. The corner encounter is violent, with both males and females flying in and out of the corner reminiscent of a fight. Eventually the females escape and race up the E wall, the males in fast pursuit. Just as they approach the NE corner, the males break to the left to chase a female that has bolted for the NW corner. As the female(s) move down the W wall, the males turn in pursuit, then turn again as the females dash for the SW corner again. The groups constantly split up and reform. Visually, the population gives the distinct appearance of deliberate activity.

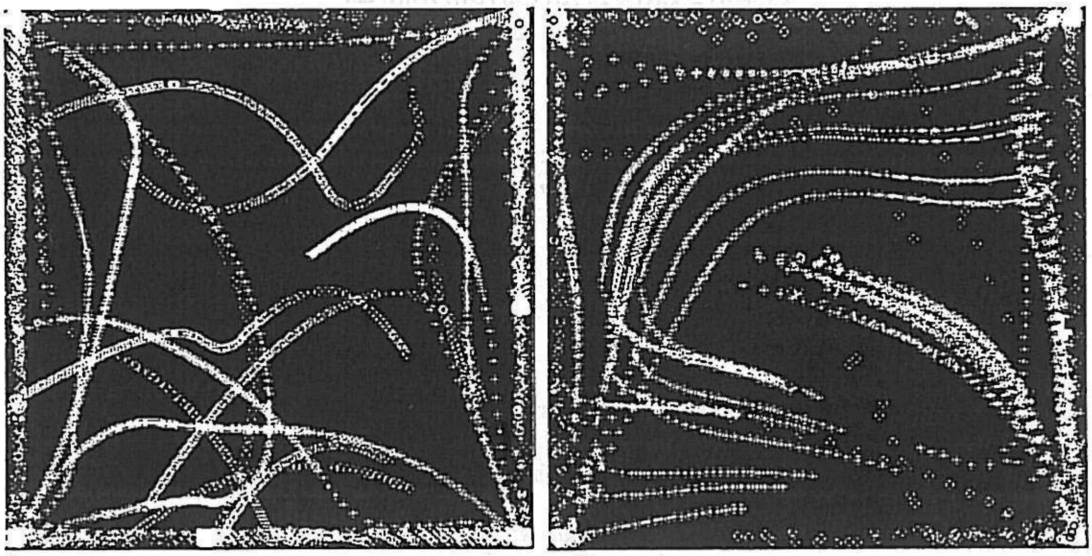

A B

Figure 4—Effects of knowledge on the collective motion. (A) Low mutual knowledge; (B) High mutual knowledge. These motions are described more completely in the text.

By adjusting the values of a, b, and S, a variety of different collective motions can be generated. For instance, setting a=1, b=0, and S={0,0.3,-0.7,0} in a population of two males and two females produces two tightly coupled male-female pairs that race around the box like scurrying mice. Infrequently, the mice collide, the individuals are scattered and mixed but quickly reform into two new mice which continue their scurrying. With a=1, b=0, and S={-1,1,1,-1}, a large population behaves like two swarms, reminiscent of gnats; the swarms drift around the box, constantly shifting positions to stay as far away from each other as possible. Still another class of behavior is found by setting $0<S_{fm}<S_{mf}$, which makes possible the creation of quasi-stable "crystals." For instance, S={0,3,1,0} in a population of 4 males and 4 females can produce a pattern with one female in each corner, three of the males stably held near the center, and the fourth male stably positioned off-center toward one edge of the box. These patterns are metastable—they are hard to produce without slowing varying S as the pattern evolves, and if perturbed, the crystals have a tendency to melt or shatter, moving randomly like a liquid thereafter.

The effects of including sex change ($b\neq0$) are spectacular. Because the motions are highly dependent on the relative sexes, a single sex change can cause the entire population to suddenly radically alter its motion, much like fish schools suddenly reversing direction. The biggest effect occurs when there is no male-male or female-female interaction: if the population is initially all the same sex, it cannot start moving, and if a moving population somehow achieves unisexuality, it stops completely and permanently.

Physical Meaning of the Algorithm

Knowledge, in the conventional use of the word, implies a test: we "know" something if we can give the correct answer to a question about that something. We imagine asking individual {i} for some statement about some property of individual {j}. If {i} has knowledge of the state of {j}, then {i} gives the correct answer; otherwise it cannot. In adapting this to the KtA, we say that if {i} has knowledge K_{ij} of {j} then {i} will have probability p_{ij} of giving the correct answer.

We specify the relationship between knowledge and probability as follows: Assume that the population is comprised of a set of identical individuals, each with a state variable that can assume any one of G possible values. If individual {i} knows nothing of individual {j}, {i} has the random probability 1/G of correctly identifying the state of {j}. If, however, {i} has knowledge K_{ij} of {j}, this probability is

$$p_{ij} = \frac{1}{G} + \left[1 - \frac{1}{G}\right] K_{ij}$$

The picture here is that {j} is in only one of the G states accessible to it, and {i} makes successive (randomly chosen) guesses of which state. {i} correctly guesses the state of {j} only a fraction p_{ij} of the time.

The property of individual {j} symbolized by the matrix element K_{ij} is quite arbitrary: it can be location, size, age, sex, color, condition, or any other property. For each property we wish to track, we specify a separate matrix K. Intrinsic to these definitions is the idea that the individuals have a specified state. K represents knowledge of the states of the individuals by all other individuals.

While the basic physical dynamics provides for *all possible* motions, the KtA will enable us to determine the *actual* motion. This is effected by introducing into the individual dynamics the probabilities that the individuals know the states of other individuals, thereby providing a weighted response of the individual to all others.

To clarify these ideas with an example, consider a collection of objects whose dynamical variables $\{x_i\}$ satisfy the equation of motion

$$\frac{dx_i}{dt} = \Sigma_j f(x_i, x_j).$$

This relation assumes that all objects in the population are simple and they have no uncertainty about the states of the other objects. If the objects are sufficiently complex that they might be uncertain about the other objects, we replace this with the equation

$$\frac{dx_i}{dt} = \Sigma_j f'(x_i, x_j, p_{ij})$$

plus the equations for p_{ij} (given in terms of K_{ij} above).

What we have done here is to factor the dynamical equations into two sets of equations, one set (as above) that weights individual behavior according to how much the individual knows about its interacting partners, and another set (the KtA) that describes the dynamics of the knowledge itself. Both sets of equations are written from physically reasonable models.

The KtA provides a hybrid dynamics specifically designed for populations of objects that have sufficient internal complexity to have some knowledge of other objects and to take actions based on that knowledge. It is most appropriately applied to simple animal behavior such as schooling or flocking, in which the full complexity of the organism is not necessary to produce the collective behavior. Whether the KtA can be meaningfully extended to populations of simple inanimate objects such as grains of sand that have no capacity for storing and processing knowledge in the conventional sense is a subtle question about physics that is as yet unanswered.

Relation to Information Theory

In the classical definition of information (Brillouin, 1962), we assume we have a system with R possible configurations. Initially, we have no knowledge which configuration the system actually has. If we subsequently acquire an amount of information $I = c \ln(R/R')$, where c is a units constant, we would infer that the system has only R' possible configurations.

In order to calculate R/R', we note that the reciprocal of p_{ij}

$$g_{ij} = G \frac{1}{1+(G-1)K_{ij}}.$$

is the number of states that {i} would infer that {j} has available to it. The total number of configurations of the population, inferred by {i}, is then $R_i = \Pi_j \, g_{ij}$. Thus, the information $I_i(t)$ associated with a transition from an initial knowledge $K_{ij}(0)$ to knowledge $K_{ij}(t)$ at time t will be

$$I_i(t) = c \ \ln\left[\frac{R_i(0)}{R_i(t)}\right] = \Sigma_j I_{ij}(t)$$

where

$$I_{ij}(t) = c \ \ln\left[\frac{1+(G-1)K_{ij}(t)}{1+(G-1)K_{ij}(0)}\right].$$

represents the self (j=i) and mutual (j≠i) information. We can invert this to obtain

$$K_{ij}(t) = \left[\frac{1}{G-1} + K_{ij}(0)\right] e^{I_{ij}(t)/c} - \frac{1}{G-1}$$

which gives the increase in knowledge K_{ij} due to information I_{ij}. The exponential relationship is a natural result of mapping an additive space into a multiplicative one.

Relation to Other Work

This work is very closely related to *harmony theory* (Smolensky, 1989), which is built on a paradigm of activation of knowledge atoms and their assembly into context-sensitive schemata. Both assembly and inference (which is the completion of missing parts) are achieved by finding maximally self-consistent states of the system. The theory leads inevitably to the definition of a harmony function H, and its exponential connection with probability: $p \sim \exp(H/kT)$. Thus, we find linear correspondence between probability and knowledge, and between information and harmony.

This work also draws strongly from the ideas of Kampis (1991), who emphasizes the limitations of pure dynamics to deal with reality. While staying within the confines of micro physics, Kampis stresses that qualitatively new behavior results if the components of a system have cognitive ability. Here, we phrase this as determining what *will* happen from what *could* happen, and we seek a prescription for writing the dynamics of such a system that displays the observed behavior.

Shoham (1993) implements a calculus for agent-oriented programming (AOP) that tracks the activities of agents capable of cognitive social functions such as beliefs, decisions, capabilities and obligations. While this work succeeds in reducing very complex social networks and chains to their irreducible effects, our purpose here is different: to simulate populations of minimally cognitive objects in real time.

Genesereth (1989) has studied collections of agents that can obtain knowledge and take actions based on partial knowledge, systems similar to the present work.

In knowledge representation theory (Harmon and King, 1985), a piece of knowledge is accompanied by a confidence factor that is the probability that the knowledge is correct. The matrix element K_{ij} can be identified as essentially this confidence factor; the KtA provides a particular dynamics for the matrix.

Many workers are producing surprisingly realistic simulations of collective behavior of real animals. Typical is the work of Huth and Wissel (1992) on fish schools, in which individual fish know only about their nearest neighbors. We believe that phrasing such simulations in terms of mutual knowledge would be useful.

In contrast, models of emergent behavior based on simple local dynamics are quite different from the dynamics discussed here. For instance, Millonas (1992) and others have modeled ant swarms based on a small set of simple rules for each ant. While this procedure does generate complex emergent behavior that often accurately simulates real ant swarms, there is no provision for nonlocal interactions that are at work in collections of individuals that know about each other.

Conclusions

In this work, I have proposed an algorithm that tracks the interactions between individuals in a population based on a metaphor of mutual knowledge. The algorithm has the effect of determining the actual motion from all possible motions. This is a hybrid dynamics, specifically designed for minimally knowledgeable individuals. It is linked through probability to mechanics, dynamics, information theory, and statistical mechanics. We find that a spectacularly small amount of knowledge is necessary for even very complex collective behavior, and that even an infinitesimal amount of knowledge is capable of producing major qualitative alterations in the evolution of the population, due to the long-term accumulation of small changes.

References

L. Brillouin. 1962. *Science and Information Theory*. 2nd Ed. Academic Press, New York.

D. L. DeAngelis, ed. 1992. *Individual-based Approaches in Ecology*. Chapman and Hall, London.

H. Gutowitz. 1991. *Cellular Automata*. Reprinted from *Physica* D **45** as special issue.

P. Harmon and D. King. 1985. *Expert Systems*. J. Wiley & Sons, New York.

A. Huth and C. Wissel. 1992. The Simulation of the Movement of Fish Schools. *J. Theor. Biol.* **156**, 365-385.

G. Kampis. 1991. *Self-Modifying Systems in Biology and Cognitive Science*. Pergamon Press, Oxford.

C. Langton, ed. 1989. *Artificial Life, Proc. Interdisciplinary Workshop on the Synthesis and Simulation of living Systems*. Los Alamos, NM, Sept. 1987, Addison-Wesley Publ. Co., Redwood City, CA.

M. M. Millonas. 1992. A Connectionist Type Model of Self-Organizing Foraging and Emergent Behavior in Ant Swarms. J. *Theor. Biol.* **159**, 529-552.

Y. Shoham. 1993. Agent-oriented programming. *Artif. Intellig.* **60**, 51-92.

P. Smolensky. 1989. Information Processing in Dynamical Systems: Foundations of Harmony Theory. in: *Parallel and Distributed Processing*, D. E. Rumelhart and J. L. McClelland, Eds., MIT Press, Cambridge.