J Mol Evol (1988) 27:114–120

Journal of
Molecular Evolution

A Model of Autocatalytic Replication

L.F. Abbott[1]

Department of Physics, Boston University, Boston, Massachusetts 02215, USA

Summary. The catalytic effects that existing polymer chains have on the formation of new chains are modeled using ideas from spin glasses and neural networks. Computer simulation shows that isolated groups of chains in this model are capable of accurately replicating a wide variety of complex structures without templating. Replication in the model arises spontaneously and rapidly, leading to an extremely simple realization of a system exhibiting Darwinian evolution.

Key words: Molecular evolution — Autocatalytic replication — Network model — Computer simulation

Introduction

In order to exhibit Darwinian evolution a system must be capable of replicating in a variety of distinguishable forms. A key difficulty in constructing either mathematical models or simple chemical systems which display Darwinian evolution is that all of these requirements—replication, variability, and distinguishability—must be satisfied simultaneously. In models based on templating (Blum 1962; Usher 1977; Kuhn and Kuhn 1978; Eigen and Schuster 1979; Bresch et al. 1980; Niesert et al. 1981; Anderson 1983; Stein and Anderson 1984; Rokhsar et al. 1986; Tsallis and Ferreira 1986), replication of a wide variety of polymer sequences is possible, but in order to distinguish between different sequences the information they contain must be expressed in characteristics that can then be selected. This requires the development of some form of genetic code. Models based on catalytic properties (Calvin 1969; Rossler 1971; Dyson 1982, 1985; Farmer et al. 1986) have expressed characteristics but need some form of autocatalytic replication to be complete. Here, I will follow the second of these strategies, considering an autocatalytic model in order to address the replication problem.

The model is based on a mathematical idealization of the behavior of self-catalyzing polymer chains. Of course, the ultimate answer to any questions about the dynamics of such chains can only come from chemistry (for reviews of the relevant chemistry see for example Oparin 1966; Fox and Dose 1977; Miller and Orgel 1985). However, mathematical modeling is useful for suggesting at least preliminary answers to questions such as these: Is it possible to construct an autocatalytic system capable of replication without templating? How simple can such a system be and what are the essential elements needed to make it work? Can it exhibit Darwinian evolution?

The model I will discuss uses the dynamics of a spin glass and ideas from neural networks to model the behavior of polymer chains in a chemically active catalytic environment. A spin glass model has been used by Anderson and collaborators (Anderson 1983; Stein and Anderson 1984; Rokhsar et al. 1986) to study the formation of nucleic acid sequences by templated polymerization. Since their work was based on templating, replication was incorporated into the structure of the model from the start. The spin glass dynamics were used to model the survival probability of polynucleotide chains and thus to select certain chains over others. It was stressed (Anderson 1983) that the spin glass analogy is partic-

Offprint requests to: L.F. Abbott
[1] On leave from Brandeis University, Waltham, Massachusetts 02254, USA

ularly useful because of the large number of stable states in the spin glass model. An approach (Stein and Anderson 1984) for modeling enzymes as well as nucleic acid chains was also discussed. However, in the computer simulations reported by Rokhsar et al. (1986), selection occurred only with respect to supposed chemical properties of nucleic acid chains, there was no possibility expressing the information contained in specific nucleic acid sequences.

Although the present work has certain similarities to that of Anderson and collaborators, it is actually more closely related to models of neural networks. Catalysis, not templating, is being considered. The spin glass model is used to simulate the dynamics of the interactions of polymer chains, not their survival probabilities. Replication is not included from the start but must instead arise spontaneously from these dynamics. The present work uses both the diversity of stable spin glass states and the ability of the model to store and recover information as in neural network models. It is this second property that allows for replication.

The model is based on the following chemical analog. Consider an isolated system containing a number of catalytically active polymer chains. Food chains in the form of high-energy polymers are introduced into the system and are transformed to lower energy forms through interactions controlled by the catalytic activity of preexisting chains. As they are transformed, the new chains (if they become catalytically active) will modify the chemical properties of the system. Food chains are introduced repeatedly until the number of catalytically active chains doubles. Then, the system is split in two, the active chains being divided up randomly. Finally, the process is repeated on the resulting daughter systems. It is clear that such a system will have achieved replication if properties of its internal chemical environment are preserved from generation to generation, that is, through successive splittings. The requirement of variability will be met if replication occurs for a wide range of these properties. It is essential that any variable characteristics be due to properties of the system and not to properties of the food chains. To assure this, I will consider the introduction of a single, unique, though randomly chosen, food chain. All variation in replicating structures is then due to internal properties and is achieved despite the lack of variability in the food supply.

In the model presented here, replication arises rapidly and spontaneously in an autocatalytic system cycled in the way discussed above. Replication is possible for a wide variety of forms and preferential selection of properties of these forms results in a remarkably simple realization of a system exhibiting Darwinian evolution.

The Model

We wish to model both the behavior of polymer chains in a chemically active environment and the way in which that environment is modified by the presence of the chains. We are only interested in the behavior of catalytically active chains and so will completely ignore inactive chains. An isolated system will be represented by n chains of length N denoted by S_i^a with $i = 1, 2, \ldots, N$ labeling sites along a chain and $a = 1, 2, \ldots, n$ labeling chains. To further simplify the model we will consider chains consisting of only two basic monomers represented by the values $+1$ and -1, so that $S_i^a = \pm 1$. The effects of the chemically active internal environment of such a system on an arbitrary chain ξ_i is determined by a matrix M_{ij} representing the effects of all the catalytically active chains S_i^a. In one interaction time (defined rather arbitrarily) the chain ξ_i, when exposed to the environment M, is transformed into a new chain ξ_i' given by

$$\xi_i' = \operatorname{sign}\left(\sum_{j=1}^{N} M_{ij}\xi_j\right) \tag{1}$$

The model makes no attempt to account for the complex series of chemical interactions that would be needed to realize this transformation. Following ideas from neural networks (Hopfield 1978), we imagine that the dynamics of the transition (1) are governed by an effective free energy

$$E = -\tfrac{1}{2}\left(\sum_{i=1}^{N}\sum_{j=1}^{N} M_{ij}\xi_i\xi_j\right) \tag{2}$$

Because of this we can impose the conditions $M_{ij} = M_{ji}$ and $M_{ii} = 0$. Then, the transformation (1) always lowers E unless it is trivial. The repeated transformation of a chain by (1) until a fixed point is reached is just the relaxation of an unstable chain to a stable configuration at a local minimum of E. The E given above is the Hamiltonian for the infinite range Ising spin glass (Sherrington and Kirkpatrick 1975). There exist a large number (of order $e^{0.2N}$) of local minima of this energy so there will be a large variety of stable chains in our model.

To complete the model we must specify how the interaction matrix M depends on the catalytically active chains present in the system, S_i^a, for $a = 1, 2, \ldots, n$. First, there must be a piece in M which is independent of these chains representing the energy of a free isolated chain. We denote this by a matrix A_{ij} introduced with strength α. The catalytic activity of the chains is represented by another term $C_{ij}(S)$, which does depend on the identity of the chains present in a particular system. Thus,

$$M_{ij} = \alpha A_{ij} + C_{ij}(S) \tag{3}$$

I have tried various forms for the matrix $C_{ij}(S)$. A particularly useful one is

$$C_{ij}(S) = \sum_{a=1}^{n} B_{ij} S_i^a S_j^a \tag{4}$$

There is nothing particularly unique about this form as far as replication is concerned, but it has the advantage of being easy and fast to implement in computer simulations and is used in all the results reported below. The fixed matrices A and B then completely determine the dynamics of the model. The matrix A is supposed to reflect the stability of various free chains and B the catalytic effect of one chain on another. Since the detailed chemistry is not known I will choose these two matrices randomly with the only restrictions being those on M,

$$A_{ij} = A_{ji} \qquad B_{ij} = B_{ji} \tag{5}$$

and

$$A_{ii} = B_{ii} = 0 \tag{6}$$

The normalization of A and B is irrelevant. In the work we discussed below these matrices are determined by randomly setting their independent and nonzero elements to $+1$ or -1 with equal probability.

The relative importance of the free chain term A and the catalytic term C in M [Eq. (3)] is determined by the parameter α. The transition to what I refer to as a chemically active system occurs when the second term in (3) is roughly comparable to the first. Then, the dynamics are governed by the catalytic interactions of the chains. Since individual terms in the sum determining C, (4), are uncorrelated, this will occur when the number of active chains is roughly of the order α^2. The parameter α^{-1} is a measure of the catalytic strength of the chains relative to typical free energy differences between chains.

The dynamics of a system of n chains of length N is determined by the updating rule (1) and the rules for forming the interaction matrix (3) and (4). The only modification needed is the observation that a chain cannot catalyze itself, so for the chains that make up a given system, S_i^a, the updating rule is

$$(S_i^a)' = \operatorname{sign}\left[\sum_{j=1}^{N} (M_{ij} - B_{ij} S_i^a S_j^a) S_j^a\right] \tag{7}$$

Eqs. (3), (4), and (7) completely specify the dynamics of the model.

The computer simulations of this model proceed as follows. First, we randomly choose the matrices A and B and the food chain which is to be added by setting their elements randomly to $+1$ or -1 with equal probability. Once these are chosen they remain the same throughout the simulation. We then begin with a group of $n = n_f/2$ chains of length N. These are also chosen randomly with $+1$ or -1 occurring with equal probability at each site S_i^a. The transformation matrix for this set of chains is computed from Eqs. (3) and (4) and all n chains are updated according to rule (7). Updating is repeated until all the chains have reached fixed points of this transformation, in other words, until all chains have reached stable configurations at local minima of E. This takes only a relatively small number of iterations. Next, a food chain is added, increasing the size of the system from n to $n + 1$ chains. All chains, including the added food chain, are updated (7), and the added food chain is now included in the determination of a new transformation matrix M using (3) and (4). Once again the transformation (7) is repeatedly applied until all chains have reached stable fixed points. Then, another food chain is added and the procedure is repeated until the number of chains has increased to $n = n_f$. At this point the system is divided randomly into two subsystems each containing $n_f/2$ chains. These are the two daughter systems and the process outlined represents one generation of evolution.

Results

Only relatively small systems have been studied thus far. (For large N the system can be studied analytically—this will be presented in a separate publication.) The results of this section refer to chains of length 10 in systems with a maximum of 20 chains. That is, $N = 10$ and $n_f = 20$. Similar results are obtained with systems in this general range of size.

To study the evolution of these systems we need a way of characterizing an individual system using some sensitive measure of the chemical properties implied by the matrix M. A useful measure is

$$t = \frac{1}{2N\sqrt{nN}} \sum_{i=1}^{N} \sum_{j=1}^{N} M_{ij} \xi_i \xi_j \tag{8}$$

where ξ_i is a randomly chosen test chain. Unless otherwise indicated, all tests are performed when the system has attained its full size $n = n_f$. Although t is a single number, it turns out to be a very useful characterization of the whole matrix M and it will be used extensively below.

To check for replication, we cycle a randomly chosen system through many generations as outlined above. We begin by considering highly cata-

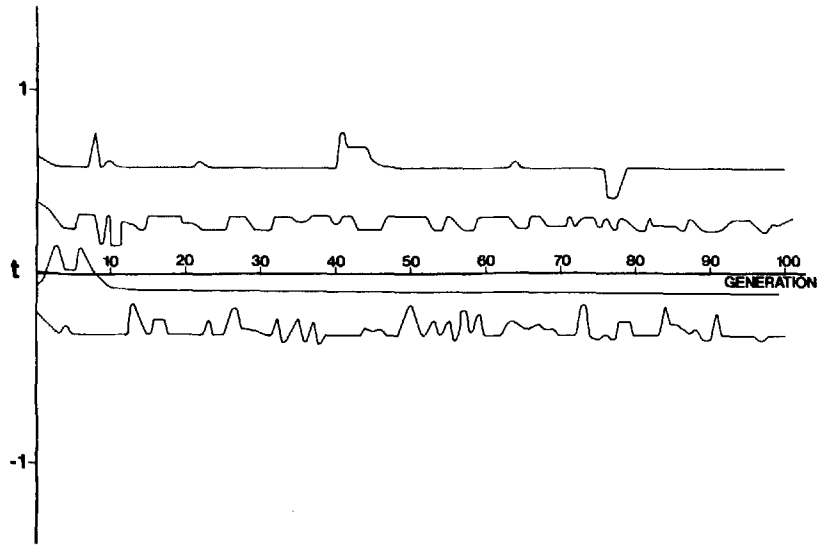

Fig. 1. The value of t as a function of generation number for four different replicating systems

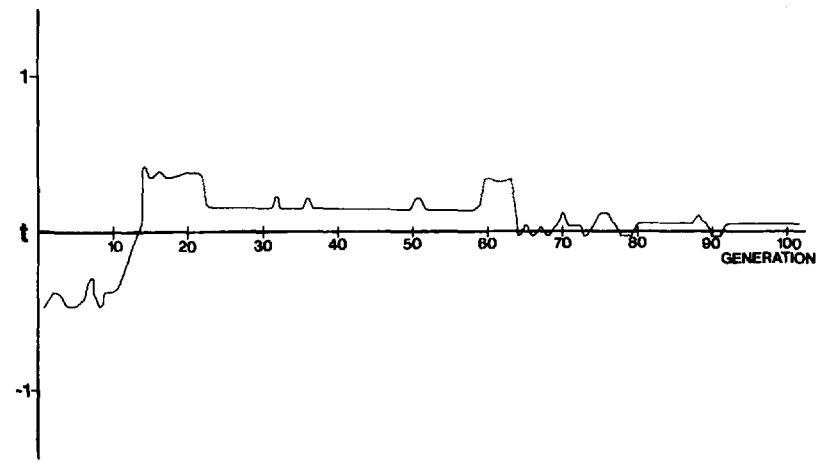

Fig. 2. Same as Fig. 1 but for a less stable system

lytic systems, so we take $\alpha = 0$ in Eq. (3). Later, we will explore the consequences of nonzero α. Once a randomly chosen starting configuration has increased in size, due to the addition of food chains, to $n = n_f$, we examine its complement of chains, measure its t value [using Eq. (8)], and divide it up randomly into two daughter systems. We then follow the evolution of one of these daughter systems, choosing which one to follow randomly. When this daughter system reaches full size, we again examine the chains that are present, measure its t value, divide it up, and follow one of its daughter systems. Repeating this procedure over many generations, we find that a substantial fraction of initial configurations develop into stable replicating structures. This is most easily seen in these computer simulations simply by examining the chains on a display screen and noting that their structure is reproduced generation after generation. However, such a form of display is impossible in a journal article. Fortunately, the variable t provides an accurate measure of replication. In all cases where the variable t remains constant or nearly constant over many generations, a detailed examination shows that exact or nearly exact replication of chain structure is taking place. Therefore we use here the variable t as a shorthand indication of the complete structure of a system. Plotting t as a function of generation number produces the curves shown in Figs. 1 and 2.

The ability of these systems to replicate is shown in Fig. 1 where the value of t for subsequent generations of daughter systems is plotted as a function of generation number for four different starting configurations. Clearly the systems are capable of maintaining a variety of individual "chemical" identities, characterized by t, over many generations. Note that one of the systems shown in Fig. 1 is extremely stable, one is very stable except for occasional fluctuations, and the other two oscillate around stable values of t. Although not all matrices A and B and not all starting configurations lead to replicating systems, it should be stressed that the systems shown in Fig. 1 are quite typical. Figure 1 also shows how varied the systems can be. Replication occurs for

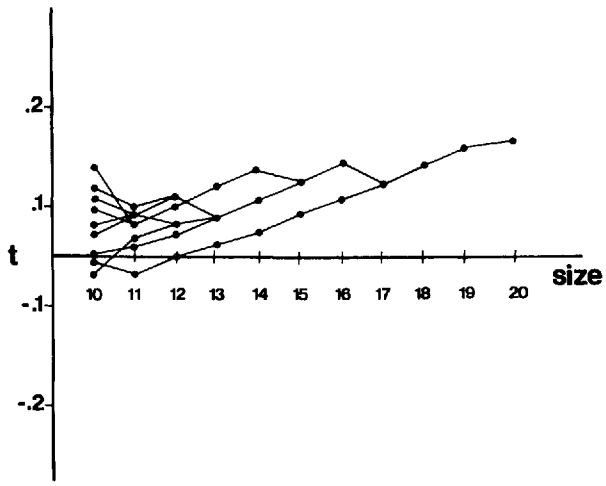

Fig. 3. The value of t as a function of size for several daughter systems all evolving to the same final configuration from different starting configurations

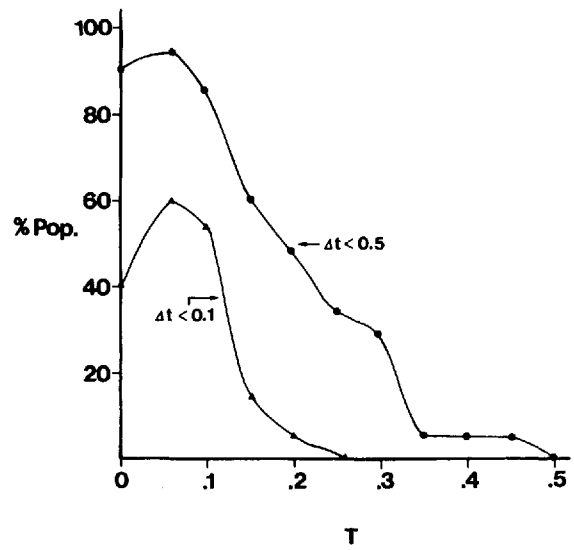

Fig. 4. The percentage of a random initial population capable of replicating with the indicated degree of accuracy as a function of temperature

virtually any value of t in the range from $t \approx -1$ to $t \approx +1$. More careful examination of these systems shows that they are indeed replicating the complete structure of their chains S_i^a, so the fact that t does not vary much over many generations is an accurate indication of true replication. Figure 2 shows a somewhat more complex evolution. This system took longer to establish a replicating configuration and then at generation 60–65 jumped from one pseudostable configuration to another due to a fluctuation in an initial daughter configuration.

Systems in this model are capable of both replication and variety because they are able to find a stable middle ground between maximally complex structures, which lead to chaotic behavior, and trivial structures, which allow replication in only a single, unique form. To show this we need a measure of the complexity of a system of chains. Suppose we have a group of n_f chains of length N, and that in this group there are m different varieties of chains (in other words, of the n_f chains there are m distinct types). Also imagine that in this group there are $N - M$ sites i at which the value of S_i^a is the same for all a. A reasonable measure of the complexity of such a group is the amount of information that can be stored in the m distinct chains using the M distinct bits. This is

$$I = m \log 2^M = mM \log 2 \qquad (9)$$

Dividing by the maximum value of I, $n_f N \log 2$, we define the relative complexity of a group of chains as $c = mM/n_f N$. For the replicating systems being discussed here, the complexity ranges from $c \approx 0.05$ to $c \approx 0.2$ so the systems are far from trivial but also far from maximally complex. A typical system might have, for example, five different types of chains varying in four different sites along the chain. For this system, c would be 0.1.

Replicating systems in this model show a remarkable degree of stability. For example, when a full-sized system consisting of 20 chains with perhaps five different types of chains is split into two daughter systems of 10 chains each, there will be significant random variations in the composition of the daughter systems. Nevertheless, the daughter systems are capable of regenerating those chains that they lack and modifying those chains that they have in excess in order to produce a final configuration of 20 chains identical to the parent configuration. This is shown in Fig. 3, where we plot the value of t for various daughter systems as they increase in size (due to the addition of food chains) from $n = 10$ to $n = n_f = 20$. The different lines indicate the various paths that the daughter systems take from the several starting configurations shown to attain identical final configurations. The ability of these systems to convert existing chains and modify food chains to achieve stable replication is particularly striking when the full computer simulation can be viewed. Of course, this is impossible here but the plots of t give an indication of what is going on.

The stability exhibited in Fig. 3 suggests that the systems might be quite tolerant of errors in the updating procedure, Eq. (7). Such errors can be simulated by running the systems at finite temperature T. This means that with probability $e^{-1/T}$ we allow a transition $S_i^a \rightarrow -S_i^a$ for a given i and a, which raises rather than lowers E, producing an error rate per site per chain per iteration of $e^{-1/T}$. In these simulations each chain is updated five times between successive additions of a food chain, so in one generation each of the original chains is updated 50 times. Figure 4 shows the percentage of random starting systems that are capable of replicating to a

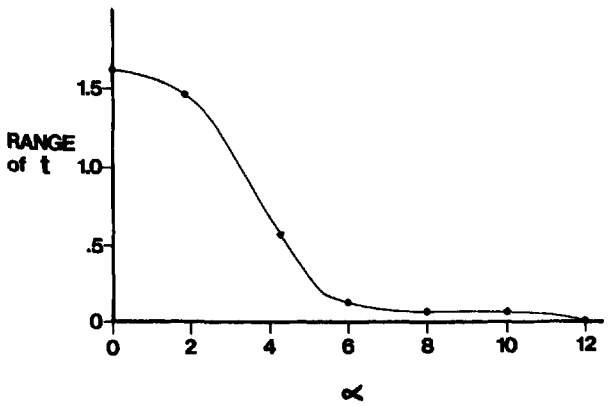

Fig. 5. The range of t values obtained for various replicating systems as a function of α

desired degree of accuracy as a function of temperature. To measure the ability of a system to replicate we use the quantity

$$\Delta t = \frac{|t - \bar{t}|}{|\bar{t}|} \tag{10}$$

averaged over five generations where $\bar{t}$ is the average value of t. The upper curve in Fig. 4 shows the percentage of starting configurations that are able to replicate with an accuracy of $\Delta t < 0.5$ after 20 generations and the lower curve corresponds to $\Delta t < 0.1$. It can be seen that replication at the level of $\Delta t < 0.5$ is not uncommon for T as high as 0.3. This corresponds to an error rate of 4% per chain per site per update or 30% per chain per update.

Changing the value of α away from $\alpha = 0$ has no effect on the ability of the systems to replicate. However, it does reduce the variability of the resulting systems. When α is too large, the systems cannot control the degradation of the added food chains, and since these have been chosen to have no variety, they will all relax to the same local minimum of E. As a result, the systems will all evolve to a unique state of n_f identical chains with complexity $c = 0$. Figure 5 shows the range of t values, $t_{\max} - t_{\min}$, obtained from a number of different starting configurations as a function of α. It can be seen how the wide range of possible t values is restricted at large α. Variability essentially stops above $\alpha = 5$.

Since these systems have all the characteristics—replication, variability, and distinguishability—needed for Darwinian evolution, it is interesting to take a population of systems and subject them to selection. Consider a population of 100 systems each with $N = 10$ and $n_f = 20$. We select for large values of the parameter t by choosing two systems out of the population randomly and removing from the population the one with smaller t. The other system (with the larger t value) is then allowed to increase in size and split, with one daughter system taking

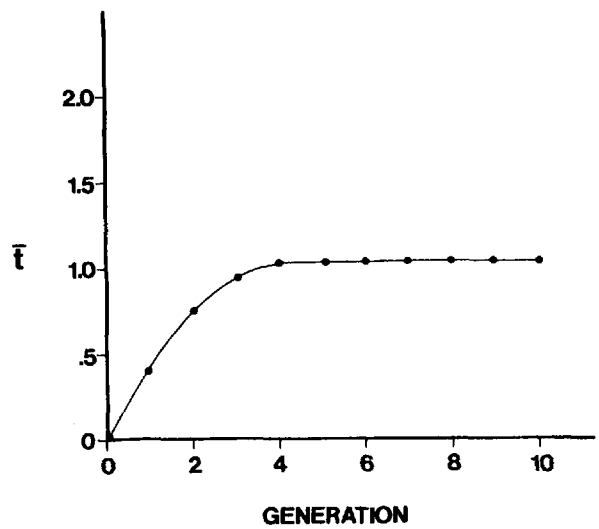

Fig. 6. The value of t averaged over an entire population selected for large t as a function of generation number

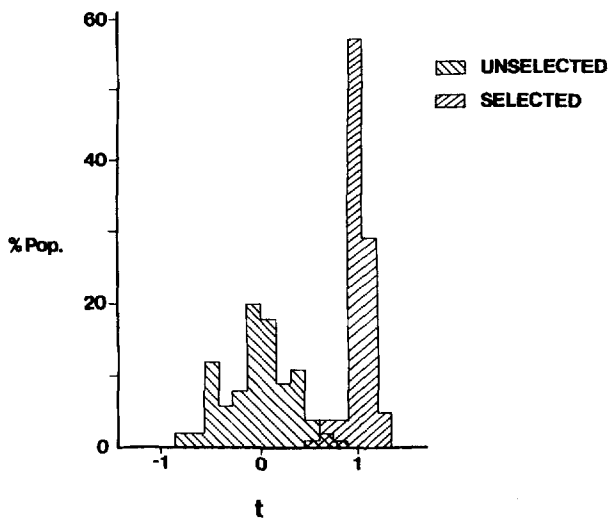

Fig. 7. The distribution of t values for an unselected and a selected population

the place of the removed element. The effect of this selection on the value of t averaged over the whole population is shown in Fig. 6. More striking is the behavior shown in Fig. 7. Here, one shaded region shows the distribution in t of an unselected population, while the other shows the t distribution for a selected population. Not only has the distribution been pushed to larger t values, but in addition the selected population is highly correlated into a narrow distribution reflecting the success of a few viable systems in dominating the final population.

Discussion

In light of the successes of the model presented here, it is interesting to consider in more general terms

the dynamical requirements for a replicating autocatalytic system. Clearly the dynamics must allow for a wide variety of well-separated stable configurations. Furthermore, in order for replication to take place, the basin of attraction (the set of unstable configurations that dynamically evolve into a given stable configuration) of a stable state must be quite broad—broad enough to include the daughter configurations. In this sense, a replicating system is a form of associative memory. In an associative memory, a suitable approximation of a given data set is sufficient to recall the entire data set accurately. In the case of a replicating system the approximate data set is just a fragment (the daughter system) of the full set (the parent system). Replication should be possible for any transformations [taking the place of (1)] with a rich set of attractive fixed points or equivalently any effective energy, E, with a rich spectrum of fairly broad stable local minima. Any form of autocatalytic action (4) strong enough to maintain the system in a minimum of E far from equilibrium should, under these conditions, be sufficient to lead to a variety of replicating structures. It remains to be seen if any nonliving chemical system meets these requirements.

Acknowledgments. This work was supported in part by Department of Energy contract DE-AC02-76ER03230 and by an Alfred P. Sloan Foundation Fellowship.

References

Anderson PW (1983) Suggested model for prebiotic evolution: the use of chaos. Proc Natl Acad Sci USA 80:3386–3390

Blum HF (1962) On the origin and evolution of living machines. Am Sci 49:474–501

Bresch C, Niesert U, Harnasch D (1980) Hypercycles, parasites and packages. J Theor Biol 85:399–405

Calvin M (1969) Chemical evolution. Oxford University Press, New York

Dyson F (1982) A model for the origin of life. J Mol Evol 18: 344–350

Dyson F (1985) Origins of life. Cambridge University Press, Cambridge UK

Eigen M, Schuster P (1979) The hypercycle. Springer, Berlin

Farmer D, Kauffman S, Packard N (1986) Autocatalytic replication of polymers. Physica 22D:50–67

Fox S, Dose K (1977) Molecular evolution and the origin of life. Freeman, San Francisco

Hopfield J (1978) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci USA 79:2554–2558

Kuhn H, Kuhn C (1978) Evolution of a genetic code simulated with a computer. Orig Life 9:137–150

Miller SM, Orgel LF (1985) The origins of life on earth. Prentice-Hall, Englewood Cliffs NJ

Niesert U, Harnasch D, Bresch C (1981) Origin of life between Scylla and Charybdis. J Mol Evol 17:348–353

Oparin AI (1966) The origin and initial development of life. Meditsina, Moscow

Rokhsar DS, Anderson PW, Stein DL (1986) Self-organization in prebiological systems: simulation of a model for the origin of genetic information. J Mol Evol 23:119–126

Rossler O (1971) A system theoretic model of biogenesis. Z Naturforsch B266:741–746

Sherrington D, Kirkpatrick S (1975) Solvable model of a spin glass. Phys Rev Lett 35:1792–1796

Stein DL, Anderson PW (1984) A model for the origin of biological catalysis. Proc Natl Acad Sci USA 81:1751–1753

Tsallis C, Ferreira R (1986) A critical phenomena approach to biogenesis. Preprint ISSN 0029-3865

Usher DA (1977) Early chemical evolution of nucleic acids: a theoretical model. Science 196:311–313

Received March 6, 1987/Revised and accepted July 23, 1987