Gérard Weisbuch

The Complex Adaptive Systems Approach to Biology

Abstract

The purpose of this paper is to describe concepts and methods inspired from statistical physics of disordered systems and non linear physics and their application to theoretical biology. The central perspective is the study of functional organization of multi-component systems, based on a simplified description of individual components. The first section discusses a few examples of complex systems in physics and biology. We then describe three basic formalisms used in theoretical biology. The most important concept of attractor is introduced in the section on networks. We will discuss generic organization properties and the difference between organized and chaotic regimes. We will then propose two possible implementations of memory in the nervous and immune systems as examples of functional organization.

Key words

Complex systems dynamics, theoretical biology, neurons, immunology, boolean nets, genetic regulatory network.

1. From Statistical Physics to Complex Systems

Statistical physics has accustomed us to mathematical descriptions of systems with a large number of components. The thermodynamic properties of ideal gases were understood as early as the end of the 19th century, while those of solids were understood at the beginning of the 20th century. In both cases, two important properties make modeling easy:

■ These are systems in which all of the components are identical.

■ If the interactions between the components are very weak, they can be ignored, as in the case of ideal gases. Otherwise, as in the case of solids, we can use linearization methods to put the problem into a form in which these simplifications can be made.

These early successes compared to the difficulties encountered in the understanding of biological systems would make us consider the above mentioned systems as rather simple. On the other hand, here are some examples of complex living systems:

■ The human brain is composed of approximately ten billion cells, called neurons. These cells interact by means of electrico–chemical signals through their synapses. Even though there may not be very many different types of neurons, they differ in the structure of their connections.

■ The immune system is also composed of approximately ten billion cells, called lymphocytes with a very large number of specificities which interact via molecular recognition, in the same way as recognition of foreign antigens.

■ Even the metabolism of a single cell is the result of the expression of a large number of genes and of the interactions among the gene products.

Although complexity is now a somewhat overused expression, it has a precise meaning within this text: *a complex system is a system composed of a large number of different interacting elements*.

In fact, the great majority of natural or artificial systems are of a complex nature, and scientists choose more often than not to work on systems simplified to a minimum number of components, which allows him or her to observe "pure" effects. The complex systems approach, on the other hand, is to simplify as much as possible the components of a system, so as to take into account their large number. This idea has emerged from a recent trend in research known by physicists as *the physics of disordered systems*.

1.1 Disordered systems

A large class of physical systems, known as multiphase systems, are disordered at the macroscopic level, but some are disordered even at the microscopic level. Glasses, for example, differ from crystals in that interatomic bonds in a glass are not dis-

tributed according to symmetries which we observe in crystals. In spite of this disorder, the macroscopic physical properties of a glass of a given composition are generally the same for different samples, as for crystals. In other words, microscopic disorder in a system does not lead to random global behavior. The simple models used by physicists are based on periodic networks, or grids, and simplified components of two different types are placed on the nodes, such as for example conductors or insulators in the problem known as percolation. These components are randomly distributed, and the interactions are limited to pairs of neighboring nodes. For large enough networks, we perceive that certain interesting properties do not depend on the particular sample created by a random selection, but of the parameters of this selection. In the case of the aforementioned insulator/conductor mixture, the conductivity between the two edges of the sample depends only on the ratio of the number of conductive sites to the number of insulating sites.

The percolation formalism exemplifies the approach taken by a number of theoretical biologists:

■ We choose to oversimplify the components of the system whose global behavior we would like to model. The formal genes, neurons and lymphocytes discussed below are cartoon-like simplifications of biological polymers and cells.

■ Nonetheless, these simplifications enable us to apply rigorous methods and to obtain exact results.

■ This approach is a dynamical approach. As in the differential methods, we start from a *local description* of the system, in terms of the short term state changes of the components as a result of their interactions. We expect the *global description* of the system from the method, that is to say the long term behavior of the system as a whole. The global behavior can be very complex, and it can be interpreted in terms of *emergent properties*. Within this notion is the idea that the properties are not *a priori* predictable from the structure of the local interactions, and that they are of biological functional significance (WEISBUCH 1990).

2. Networks

2.1 Units

Boolean automata. A simplified automaton is defined by its sets of inputs and outputs and by the *transition function*, which gives the output at time $t + 1$ as a function of the inputs and sometimes also the internal state (i.e., the output) at time t.

Boolean automata operate on binary variables, that is to say variables which take the values 0 or 1. In logical terms, 0 and 1 correspond to FALSE and TRUE, respectively. The usual logic functions AND, OR, and XOR are examples of transition functions of boolean automata with two inputs. A boolean automaton with k inputs, or of *connectivity* k, is defined by a truth table which gives the output state for each one of the 2^k possible inputs. There are 2^{2^k} different truth tables, and then 2^{2^k} automata.

Let $k = 2$. Here are the truth tables of four boolean logic functions with two inputs:

	AND	OR	XOR	EQU
Input	00 01 10 11	00 01 10 11	00 01 10 11	00 01 10 11
Output	0 : 0 : 0 : 1	0 : 1 : 1 : 1	0 : 1 : 1 : 1	1 : 0 : 0 : 1

On the input line of the table, we have represented the four possible input states by 00, 01, 10, and 11. The four truth tables correspond to the standard definitions of the following logic functions: AND returns a 1 only if its two inputs are 1; OR returns a 1 only if at least one of its inputs is a 1; XOR is 1 only if exactly one of its inputs is a 1; and EQU the complement of XOR returns 1 when both input are equal. In logical terms, if A and B are two propositions, the proposition (A AND B) is true only if A and B are true.

We will further discuss the application of boolean units to genetics.

Threshold automata. The state x_i of the ith threshold automaton is computed according to:

$$h_i = \sum_j J_{ij} x_j \,, \qquad (2.1)$$

$$x_i = 1 \text{ if } h_i > \theta_i \,;\; x_i = 0 \text{ otherwise}$$

The sum is computed over all of the inputs, subscripted by j. J_{ij} is the weight of the interaction between the ith and jth automata. In other words, the ith automaton has the value 1 if the weighted sum of the states of the input automata $\sum J_{ij} x_j$ is greater than or equal to the threshold, and 0 otherwise. Threshold automata are Boolean, but not all Boolean automata are threshold automata. We will further summarize some applications of threshold units to cognition (HERTZ/KROGH/PALMER 1990).

Formal lymphocytes. Not all networks are made of automata. A number of authors studying neural nets used differential equations as units. In immu-

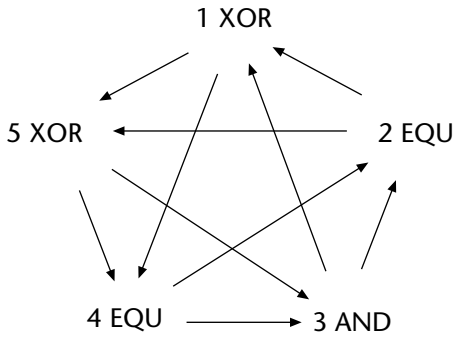

Figure 1: A network of five boolean automata with two inputs. Each automaton has two inputs and transmits its output signal to two other automata. The XOR and AND functions have been previously defined. The EQU(ivalence) function is the complement of the XOR function—it is 0 only if exactly one input is a 1.

nology, PERELSON/WEISBUCH (1997) started from the following model, called the *B model* since it deals with B cells dynamics. The time evolution of the population x_i of clone i is described by the following differential equation:

$$\frac{dx_i}{dt} = m + x_i(pf(h_i) - d), \qquad (2.2)$$

where m is a source term corresponding to newly generated cells coming into the system from the bone marrow, the function $pf(h_i)$ defines the rate of cell proliferation as a function of the "field" h_i, and d specifies the per capita rate of cell death. For each clone i, the total amount of stimulation h_i is considered to be a linear combination of the populations of other interacting clones j. This linear combination is called the field, h_i, acting on clone x_i, i.e.,

$$h_i = \sum_j J_{ij} x_j \qquad (2.3)$$

where J_{ij} specifies the interaction strength (or affinity) between clones x_i and x_j. The choice of a J matrix defines the topology of the network. Typically J_{ij} values are chosen as 0 and 1. The most crucial feature of this model is the shape of the activation function $f(h_i)$, which is taken to be a log bell-shaped dose-response function

$$f(h_i) = \frac{h_i}{\theta_1 + h_i}\left(1 - \frac{h_i}{\theta_2 + h_i}\right) = \frac{h_i}{\theta_1 + h_i}\frac{\theta_2}{\theta_2 + h_i}, \qquad (2.4)$$

with parameters θ_1 and θ_2 chosen such that $\theta_1 << \theta_2$. Below the maximum of $f(h_i)$, increasing h_i increases $f(h_i)$; we call this the *stimulatory regime*. Above the maximum, increasing h_i decreases $f(h_i)$; we call this the *suppressive regime*. When plotted as a function of log h_i, the graph of $f(h_i)$ is a bell-shaped curve.

2.2 Structural Properties

An *network* is composed of a set of units interconnected such that the outputs of some are the inputs of others. It is therefore a directed graph, where the nodes are the units and the edges are the connections from the output of one unit to the input of another. Figure 1 represents the graph of the connections of a network of five boolean automata with two inputs. This graph is equivalent to a set of five logical relations:

$$e(1) = XOR\ (e(2), e(3))$$
$$e(2) = EQU\ (e(3), e(4))$$
$$e(3) = AND\ (e(4), e(5))$$
$$e(4) = EQU\ (e(5), e(1))$$
$$e(5) = XOR\ (e(1), e(2))$$

where $e(i)$ is the state of the ith automaton.

2.3 Dynamical properties

Iteration mode. The dynamics of an automata network are completely defined by its connection graph (which automaton is connected to which), the transition functions of the automata, and by the choice of an *iteration mode*: It must be stated whether the automata change their state simultaneously or sequentially, and in what order. In the parallel mode, for instance, all of the automata change their state simultaneously as a function of the states of the input automata in the previous time step. Conversely, in the case of *sequential iteration*, or iteration in series, only one automaton at a time changes its state. Sequential iteration is therefore defined by the order in which the automata are to be updated. In the discussion that follows, we will talk only of *parallel iteration*.

Iteration graph. There are 2^N possible configurations for a network of N boolean automata. The network goes from one configuration to the next by applying the state change rule to each automaton. Its dynamics can be represented by a directed graph, the *iteration graph*, where the nodes are the configurations of the network and the directed edges indicate the direction of the transitions of the network from its configuration at time t to a new configuration at time $t + 1$. Figure 2 represents

the iteration graph of the previous network (Figure 1) for the case of parallel iteration. This graph contains the $2^5 = 32$ possible states. It illustrates the fundamental dynamical characteristics which we will define below.

Attractors. Since an automata network is a deterministic system, if the network reaches a state for the second time, it will go through the same sequence of states after the second time as it did after the first time. Therefore, the system will go into an infinite loop in state space. These loops are called the *attractors* of the dynamical system, and the time it takes to go around the loop is called the *period* of the attractor. If this period is 1, as is the case for the configuration numbered 8 in the example, the attractor is a *fixed point*. We speak of *a limit cycle* if the period is greater than 1. The set of configurations which converge toward an attractor constitutes a *basin of attraction*. The network shown in the example below has four attractors.

Clearly it is only possible to construct a complete iteration graph for small networks. For the large networks we must be content to describe the dynamics of the system by characterizing its attractors. In this way we can try to determine:

- The number of different attractors,
- Their periods,
- The sizes of the basins of attraction (the number of configurations which converge toward each attractor),
- The notion of *distance* is also very important. The *Hamming distance* between any two configurations is the number of automata which are in different states.

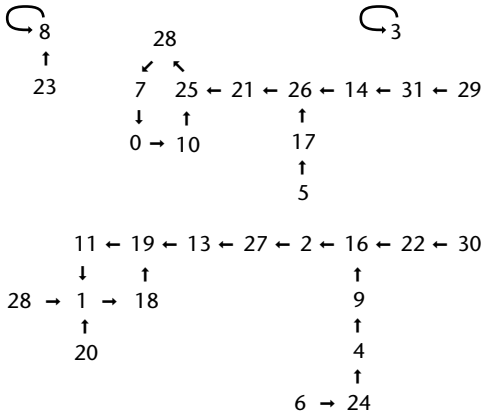

Figure 2: Iteration graph of the network of Figure 1. The numbers from 0 to 31 refer to the decimal representations of the 32 binary configurations of the network. The arrows show the temporal order of the configurations. Note that there are four different basins of attraction. State number 3 is an isolated fixed point. State number 8 is another fixed point. The other, larger, basins are composed of the configurations which converge toward the limit cycles with periods 4 and 5.

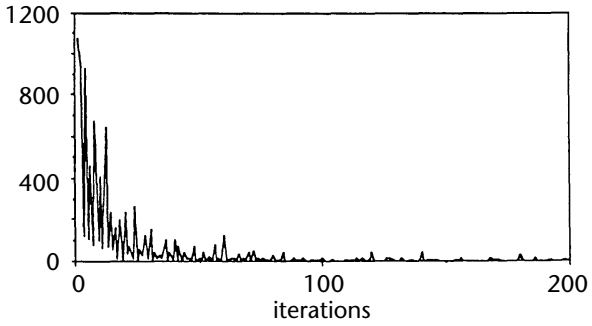

Figure 3: Histogram of the periods for 10 initial conditions of 1000 random boolean networks of 256 automata.

3. In Search of Generic Properties

In view of all the simplifications that were made to define the units of the model networks, one cannot expect all properties of living systems to be modeled. Only some very general properties, independent of the details of the model will show-up. These are the so-called *generic* properties of the network. In fact, we are interested not in the particularities of a specific network, but in the orders of magnitude which we expect to observe in studying a set of networks with fixed construction principles. We therefore consider a set containing a large but finite number of networks. We choose some of these networks at random, construct them, and measure their dynamical properties. We then take the average of these properties, and we examine those which are fairly evenly distributed over the set of networks. An example will help to clarify these ideas.

Consider the boolean networks with connectivity $k = 2$, with a random connection structure. The dynamical variable we are interested in is the period, for the set of all initial conditions and networks. Of course, this period varies from one network to the next. We have measured it for 10 randomly chosen initial conditions for 1000 different networks of 256 randomly connected automata, whose state change functions were generated at random at each node of the network. Figure 3 shows the histogram of the measured periods. This histogram reveals that *the order of magnitude of the period is ten* (this is the generic property), even though the distribution of the periods is quite large.

We can certainly construct special "extreme" networks for which the period cannot be observed before a million iterations. For this, we need only take networks which contain a random mixture of exclusive OR and EQUivalence functions (EQU is the complementary function of XOR; its output is 1 only if its two inputs are equal). But these extreme cases are observed only for a tiny fraction $(1/7^{256})$ of the set under consideration. We consider them to be pathological cases, *i.e.,* not representative of the set being studied.

We then call *generic properties* of a set of networks those properties which are independent of the detailed structure of the network—they are characteristic of almost all of the networks of the set. This notion then applies to randomly constructed networks. The generic properties can be shown not to hold for a few pathological cases which represent a proportion of the set which quickly approaches 0 as the size of the network is increased. In general the generic properties are either:

- Qualitative properties with probabilities of being true that are close to 1; or
- Semi-qualitative properties, such as the scaling laws which relate the dynamical properties to the number of automata.

The notion of generic properties characteristic of randomly constructed networks is the basis for the theoretical biological models. It has been extensively developed by physicists of disordered systems for the study of random microscopic systems such as glasses, or macroscopic multiphase systems. Physicists discovered (or rediscovered) many new concepts during the 70s. The notion of generic properties is similar to the notion of universality classes, developed for phase transitions. Without going into too much detail, we can say that the physical variables involved in phase transitions obey scaling laws which can be independent of the transition under consideration (such as, for example, phase transitions in magnetism, superconductivity, or physical chemistry) and of the details of the mathematical model which was chosen. These laws only depend on the physical dimension of the space in which the transition takes place (for us, this is three-dimensional space) and on the dimension of the order parameter. The set of phase transitions (and their mathematical models) which obey the same scaling laws constitutes a universality class. In fact, the first attempt to model a biological system by a disordered network of automata by S. KAUFFMAN (1969, 1993), a theoretical biologist, predates the interest of physicists in this subject. It is also based on the idea that the properties of disordered systems are representative of the vast majority of systems defined by a common average structure.

3.1 An example: Cell differentiation and random Boolean automata

The apparent paradox of cell differentiation is the following:

"Since all cells contain the same genetic information, how can there exist cells of different types within a single multicellular organism?"

Indeed, our body contains cells with very different morphologies and biological functions: neurons, liver cells, red blood cells… a total of more than 200 different cell types. Yet the chromosomes, which carry the genetic information, are not different in different cells. Part of the answer is that not all of the proteins coded for by the genome are expressed (synthesized with a non-zero concentration) in a cell of a given type. Hemoglobin is found only in red blood cells, neurotransmitters and their receptors only appear in neurons, etc.

Several mechanisms can interfere with the different stages of gene expression to facilitate or block it. We speak of activation and repression. The best known mechanisms involve the first steps of transcription. In order to transcribe the DNA, a specific protein, DNA polymerase, must be able to bind to a region of the chain, called the promoter region, which precedes the coded part of the macromolecule. Now, this promoter can be partially covered by a control protein, called the repressor; reading downstream gene is then impossible. It follows that, depending on the quantity of repressor present, the gene is either expressed or not expressed. The protein which acts as a repressor is also coded for by another gene, which is itself under the control of one or several proteins. It is tempting to model the network of these interdependent interactions by an automata network.

- A gene is then represented by an automaton whose binary state indicates whether or not it is expressed. If the gene is in state 1, it is expressed and the protein is present in large concentrations in the cell. It is therefore liable to control the expression of other genes.
- The action of control proteins on this gene is represented by a boolean function whose inputs are the genes which code for the proteins controlling its expression.
- The genome itself is represented by a network of boolean automata which represents the interactions between the genes.

In such a network, the only configurations which remain after several iteration cycles are the attractors of the dynamics, which are fixed points or limit cycles. These configurations can be interpreted in terms of cell types: a configuration corresponds to the presence of certain proteins, and consequently to the biological function of a cell and its morphology. Consequently, *if* we know the set of control mechanisms of each of the genes of an organism, we can predict the cell types. In fact, this is never the case, even for the simplest organisms. Without knowing the complete diagram of the interactions, S. KAUFFMAN (1969) set out to uncover the generic properties common to all genomes by representing them by random boolean networks. Since there is a finite number of possible boolean laws for an automaton with a given input connectivity k, it is possible to construct a random network with a given connectivity.

S. KAUFFMAN determined the scaling laws relating the average period of the limit cycles and the number of different limit cycles to N, the number of automata in the network. For a connectivity of 2, these two quantities seem to depend on the square root of N (although the fluctuations are very large). In fact, these same scaling laws have been observed for the time between cell divisions and for the number of cell types as a function of the number of genes per cell.

It is clear that KAUFFMAN's approximations were extremely crude compared to the biological reality—binary variables representing protein concentrations, boolean (and thus discrete) functions, simultaneity of the transitions of automata, random structures… The robustness of the results obtained with respect to the possible modifications of the model (these are random networks) justifies this approach. As for the existence of a large number of attractors, it is certainly not related to the particular specifications of the chosen networks; it is a generic property of complex systems, which appears as soon as frustrations exist in the network of the interactions between the elements. Presently, with the availability of many expression patterns, transcriptomes available thanks to DNA chips (BROWN/BOTSTEIN 1999), theoretists are facing a new challenge: how to deduce the network of gene expression regulation from the observation of the transcriptomes.

Presently, with the availability of many expression patterns, transcriptomes avail-able thanks to DNA chips (BROWN/BOTSTEIN 1999), theoretists are facing a new challenge: how to deduce the network of gene expression regulation from the observation of the transcriptomes.

3.2 Generic properties of random Boolean nets

In fact, the results obtained by KAUFFMAN show two distinct dynamical regimes, depending on the connectivity. For networks of connectivity 2, the average period is proportional to the square root of N, the number of automata. The same is true of the number of attractors. In other words, among the 2^N configurations which are *a priori* possible for the network, the dynamics selects only a small number of the order of N which are really accessible to the system after the transient period. This selection can be interpreted to be an *organization* property of the network. As the connectivity is increased, the period increases much faster with the number of automata; as soon as the connectivity reaches 3, the period as well as the number of attractors become exponential in the number of automata. These periods, which are very large as soon as the number of automata is greater than one hundred, are no longer observable, and are reminiscent of the chaotic behavior of continuous aperiodic systems. In contrast with the organized regime, the space of accessible states remains large, even in the limit of long times. Further research (see DERRIDA 1987) has shown that other dynamical properties of these discrete systems resemble those of continuous chaotic systems, and so we will refer to the behavior characterized by long periods as *chaotic*.

Functional structuring. We have shown that when boolean automata are randomly displayed on a grid their temporal organization in period is related to a spatial organization in isolated islands of oscillating automata as soon as the attractor is reached. In the organized regime, percolating structures of stable units isolate the oscillating islands. In the chaotic regime the inverse is true: few stable units are isolated by a percolating set of oscillating units (WEISBUCH 1990).

The phase transition. The connectivity parameter is an integer. It is interesting to introduce a continuous parameter in order to study the transition between the two regimes: the organized regime for short periods, and the chaotic regime corresponding to long periods. B. DERRIDA and D. STAUFFER (1986) suggested the study of square networks of boolean automata with four inputs. The continuous parameter p is the probability that the output of the automaton is 1 for a given input configuration. In other words, the networks are constructed as follows. We determine the truth table of each automa-

ton by a random choice of outputs, with a probability p of the outputs being 1. If $p = 0$, all of the automata are invariant and all of the outputs are 0; if $p = 1$, all of the automata are invariant and all of the outputs are 1. Of course the interesting values of p are the intermediate values. If $p = 0.5$, the random process described above evenly distributes all of the boolean functions with four inputs; we therefore expect the chaotic behavior predicted by KAUFFMAN. On the other hand, for values of p near zero, we expect a few automata to oscillate between attractive configurations composed mainly of 0's, corresponding to an organized behavior. Somewhere between these extreme behaviors, there must be a change of regimes. The critical value of p is 0.28. For smaller values, we observe small periods proportional to a power of the number of automata in the network. For $p > 0.28$, the period grows exponentially with the number of automata.

Distances. The distance method has recently been found to be one of the most fruitful techniques for determining the dynamics of a network. Recall that the Hamming distance between two configurations is the number of automata in different states. This distance is zero if the two configurations are identical, and equal to the number of automata if the configurations are complementary. We obtain the relative distance by dividing the Hamming distance by the number of automata.

The idea of the distance method is the following: we choose two initial conditions separated by a certain distance, and we follow the evolution in time of this distance. The quantity most often studied is the average of the asymptotic distance, measured in the limit as time goes to infinity. We compute this average over a large number of networks and of initial conditions, for a fixed initial distance. Depending on the initial distance, the two configurations can either evolve toward the same fixed point (in which case the distance goes to zero), or toward two different attractors, or they could even stay a fixed distance apart (in the case of a single periodic attractor), regardless of whether the period is long or short. Again, we observe a difference in the behaviors of the two regimes. On Figure 4, obtained with a cellular connectivity (the network is a regular two-dimensional grid), the x-axis is the average of the relative distances between the initial configurations, and the y-axis is the average of the relative distances in the limit as time goes to infinity. In the chaotic regime, we observe that if the initial distance is different from 0, the final distance is greater than 10%. The final distance seems almost independent of the initial distance. On the other hand, in the organized regime, the final distance is proportional to the initial distance.

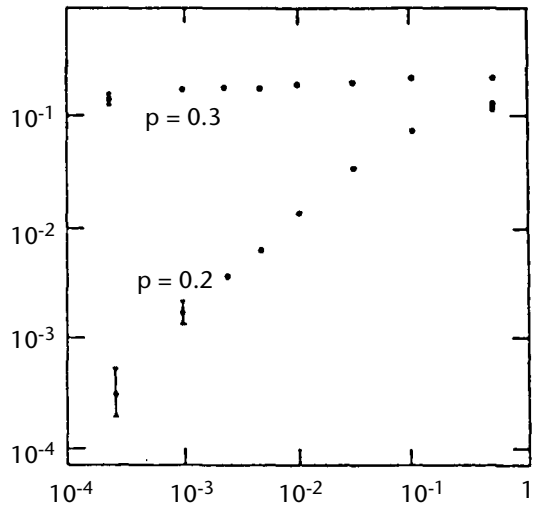

Figure 4: Relative distances at long times as a function of the initial relative distances, in the organized ($p = 0.2$) and chaotic ($p = 0.3$) regimes (from DERRIDA/STAUFFER 1986).

Property	Organized regime	Chaotic regime
Period	small	large
Scaling law (periods)	goes as a root of N	exponential in N
Oscillating nodes	isolated subnetworks	percolate
Distance	d_∞ proportional to d_0	d_∞ remains finite

Table 1: Generic properties of random networks differ according to the dynamical regime, organized or chaotic. N is the number of automata in the network, d_0 is the distance (always taken small with respect to N) between two initial configurations, and d_∞ the distance between the two evolved configurations at large times.

Conclusions. This study clearly demonstrates the existence of two types of behaviors, organized and chaotic. Table 1 summarizes the differences in the generic properties of these two regimes.

4. Memories

4.1 Neural nets and distributed memories

There now exists a very large literature on neural nets which we are not going to report here (HERTZ/KROGH/PALMER 1990). Let simply summarize the re-

sults. Neural nets with symmetrical connections have an exponential number of point attractors. This result applies to random serial iteration, and exponential means that the logarithm of number of attractors is proportional to the number of units.

Neural nets are most often used in learning tasks. A general learning algorithm is HEBB's rule. When reference patterns (network configurations) are presented to a network to be learned, connections can be constructed that ensure that the attractors of the network dynamics are the reference patterns. Furthermore the dynamics drives the network from initial conditions not to far from the reference patterns to the nearest reference patterns: these nets can then be used as associative memories that can be recalled from partial memories.

HEBB's rule can be written:

$$J_{ij} = \sum_{\mu} S_i^{\mu} S_j^{\mu}$$

where μ refers to the different reference patterns and S_i and S_j to the states of connected neurons i and j in the corresponding pattern.

Memories are thus distributed in the network as opposed to a memory that would be localized on some part of the net. The memory capacity of a fully connected neural net build according to HEBB's rule scales as the number of units in the net: no more than 0.14 N patterns, where N is the number of units, can be stored and retrieved in a HOPFIELD neural net.

4.2 Immune nets and localized memories

As a memory device, the immune system needs to obey certain constraints (PERELSON/WEISBUCH 1997): it should be sensitive enough to change attractor under the influence of antigen. It should not be too sensitive and over react when antigen is present at very low doses. The immune system should also discriminate between self-antigens and foreign antigens. Finally, it should be robust—memories of previously presented antigens should not be lost when a new antigen is presented. Thus, in some sense, the system should be able to generate independent responses to many different antigens. This independence property is achieved when attractors are localized, i.e., when the perturbation induced by an encounter with an antigen remains localized among the clones that are close to those that actually recognize the antigen (see Figure 5).

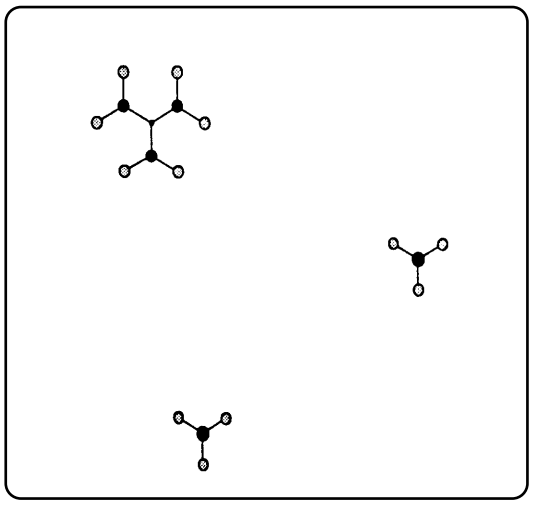

Figure 5: Localized patches of clones perturbed by different antigenic presentations. Two vaccination and one tolerant attractors are represented.

Our problem is to classify the different attractors of the network and to interpret the transitions from one attractor to another under the influence of antigen perturbation.

Let us start with the most simple virgin configuration, corresponding to the hypothetical case where no antigen has yet been encountered and all populations are at level m/d, i.e., all proliferation functions are 0. After presentation of the first antigen, memorization is obtained if some populations of the network reach stable populations different from m/d. In the case of a localized response, there will be a patch close to the antigen specific clone in which cells are excited out of the virgin state. Each antigen presented to the network will result in a patch of clones that are modified by the presentation. As long as the patches corresponding to different clones do not overlap, the various antigens presented to the network can all be remembered. Once the idea of localized non-interacting attractors is accepted, everything is simplified: instead of solving 10^8 equations, we only have to solve a small set of equations for those neighboring clones with large populations, supposing that those further clones that do not belong to the set have populations m/d. A practical approach to studying localized attractors is to combine computer simulations and analytic checks of the attractors by solving the field equations (see below).

Immunity. Let us examine the case of antigen presented to clone Ab_1, which results in excitation of

clones AB_2, clones AB_3 remaining close to their virgin level (see Figure 6). We expect that AB_1 will experience a low field, L, while AB_2 will experience a large suppressive field, H. From the field equations we can compute the populations x_i. Recall, from Eqs. (2.2) to (2.4),

$$h_1 = zx_2 = L = \frac{d\theta_1}{p'} \tag{4.1}$$

$$h_2 = x_1 + (z-1)\frac{m}{d} = H = \frac{p'\theta_2}{d} \tag{4.2}$$

where $p' = p - d$.

An immune attractor is usually reached for an intermediate initial antigen concentration, and intermediate decay constants, If the initial antigen concentration is too low or if the antigen decays too fast, the immune attractor is not attained and the system returns to the virgin configuration, i.e., AB_1 and AB_2 populations increase only transiently and ultimately return to the virgin m/d level. Thus, no memory of antigen encounter is retained.

Tolerance. Another localized attractor corresponds to tolerance (see Figure 7).

A strong suppressive field acts on AB_1 due to AB_2's, the AB_2's proliferate due to a low field provided by AB_3's, but AB_4's remain nearly virgin. The field equations once more allow one to compute the populations:

$$h_2 = x_1 + (z-1)x_3 = L = \frac{d\theta_1}{p'} \tag{4.3}$$

which gives x_3 if one neglects x_1, which is small.

$$h_3 = x_2 + \frac{(z-1)m}{d} = H = \frac{p'\theta_2}{d}, \tag{4.4}$$

and thus for small m/d

$$h_1 = zx_2 \approx zH \tag{4.5}$$

Substituting h_1 in Eq. (2.2) gives a very small value for $f(h_1)$, which shows that x_1 is of the order of m/d. The AB_1 population, experiencing a field several times higher than H, is said to be *over-suppressed*.

As in the case of the immune attractor, one can study the conditions under which the tolerant attractor is reached when antigen is presented. One finds that tolerance is obtained for large initial antigen concentrations, slow antigen decay rates and large connectivity, z (PERELSON/WEISBUCH 1997).

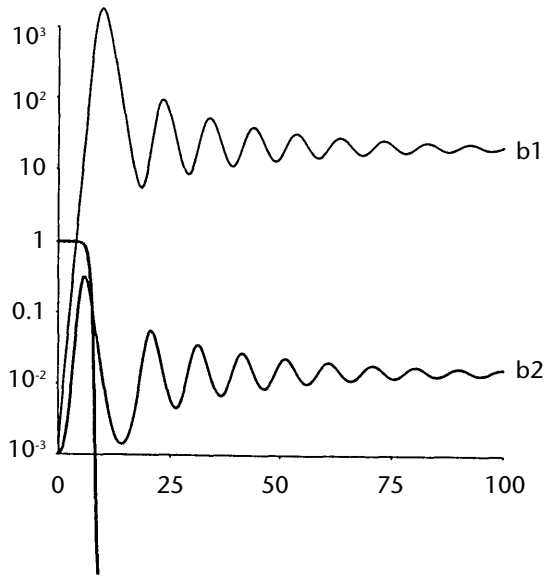

Figure 6: Time plot of an antigen presentation resulting in a vaccination attractor. On the vertical axis are the clone populations on a logarithmic scale. Time in days is on the horizontal axis. In the vaccinated configuration the largest population is localized at the first level. X_1 is high (H) and sustained by an intermediate population (L/z) of X_2. The rest of the clones are virgin (V) (or almost virgin) after the system settles into this attractor. When antigen is presented again, it is eliminated faster than the first time.

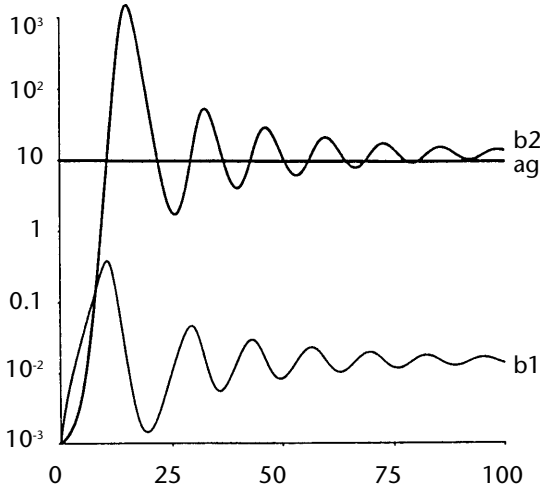

Figure 7: Time plot of an antigen presentation resulting in a tolerant attractor. X_2 is high (H) and sustained by an intermediate population (L/z) of X_3. X_1 is over-suppressed by the X_2 and is not able to remove the antigen.

4.3 Number of attractors

Localized attractors can be interpreted in terms of immunity or tolerance. Because these attractors are localized they are somehow independent: starting from a fully virgin configuration, one can imagine successive antigen encounters that leave footprints on the network by creating non-virgin patches, each of these involving a set of p perturbed neighboring clones. An immune patch contains $1 + z$ clones, a tolerant patch $1 + z^2$ (see Figure 5). Independence of localized attractors implies a maximum number of attractor configurations that scales exponentially with N, the total number of clones. The following simplified argument gives a lower bound. Divide the network into $N/(1 + z^2)$ spots. Each spot can be in 3 possible configurations: virgin, immune or tolerant. This gives a number of attractors that scales as $3N/(1 + z^2)$. Few of these attractors are of interest. The relevant question is the following: A living system must face frequent encounters with antigen during its life. Self antigen should elicit a tolerant response; dangerous external antigens should elicit immune responses and subsequent immunity. The nature of the localized response on each individual site of the network is then determined by the fact that the presented antigen should be tolerated or fought against. In this context, we can ask how many different antigens can be presented so that no overlap among different patches occurs.

In the case of random antigen presentation, simple reasoning (WEISBUCH 1990; WEISBUCH/OPREA 1994) is sufficient to derive the scaling law relating m, the memory capacity (i.e., the maximum number of remembered antigens) to N, the total number of clones. Let n_s be the number of suppressed clones involved in a patch.

m is given by:

$$m \cong \sqrt{\frac{2N}{n_s}}$$

and this provides an estimate for the mean memory capacity of the network.

The only assumption to obtain this scaling law is the random character of the network with respect to antigens, i.e., the network is not organized to respond to the set of presented antigens. On the other hand, it can be argued that the clones expressed by mammals have been selected by evolution according to the environment of the immune system, e.g., to be tolerant to self molecules and responsive to frequently encountered parasites and pathogens. If the system were optimized to the antigens in its environment, the network could be filled compactly with non-overlapping patches. The number of antigens (patches) would then scale linearly, i.e.,

$$m \propto \frac{N}{n_s}.$$

WEISBUCH/OPREA (1994) discuss more thoroughly the capacity limits of model immune networks with localized responses. They verify by numerical simulations the square root scaling law for the memory capacity. They also examine a number of other features of the network. They show that when the number of presented antigens increases, failures to remove the antigen occur since the relevant clone has been suppressed by a previous antigen presentation. They also show that previous immune or tolerant attractors are rather robust in the sense that destruction of these local attractors by new encounters with antigen is rare, and that the complete re-shuffling of the attractors, as in HOPFIELD nets (HERTZ/KROGH/PALMER 1990), is never observed.

5. Conclusions

This presentation does not mention a number of interesting topics such as the origin of Life and Species, issues in population ecology, metabolic networks, etc. The selected topics that I developed reflect only my own research interests. Since the first version of the paper was written new experimental techniques and the availability of huge series of data changes the outlook for theoretical biology.

- A number of empirical studies concerned the topology of interactions and pointed to the importance of scale free networks, i.e., networks with a wide distribution of connectivities, such as foodwebs, gene regulatory nets, metabolic networks... (ALBERT/BARABASI 2002).
- The availability of huge data sets allow to envision the solution of inverse problems: how to compute the set of interactions from the set of observed network configurations.
- But the importance of characterizing the generic properties of empirical or simulated networks still remains a preliminary step before further studies.

Author's address

Gérard Weisbuch, Laboratoire de Physique Statistique, Ecole Normale Superieure, 24 rue Lhomond, F-75231 Paris Cedex 05, France. Email: weisbuch@lps.ens.fr

References

Albert, R./Barabasi, A. L. (2002) Statistical mechanics of complex networks. Reviews of Modern Physics 74:47. lanl archive: cond-mat/0106096.

Brown, P. O./Botstein D. (1999) Exploring the New World of the genome with DNA microarrays. Nature Genetics 21:33–37.

Derrida, B. (1987) Dynamical phase transitions in random networks of automata. In: Souletie, J./Vannimenus, J./Stora, R. (eds) Chance and Matter. North Holland: Amsterdam.

Derrida, B. /Stauffer, D. (1986) Phase transitions in two-dimensional Kauffman cell automata. Europhysics Letters 2:739.

Hertz, J. A./Krogh, A. /Palmer, R. (1990) Introduction to the theory of neural computation. Santa-Fe Institute Studies in the Sciences of Complexity. Addison-Wesley: Redwood City CA.

Kauffman, S. (1969) Metabolic stability and epigenesis in randomly constructed genetic nets. Journal of Theoretical Biology 22:437–467.

Kauffman, S. (1993) The origins of order: Self-organization and selection in evolution, Oxford University Press: New York.

Perelson, A./Weisbuch, G. (1997) Immunology for physicists. Reviews of Modern Physics 69:1219–1267.

Weisbuch, G. (1990) Complex systems dynamics. Santa-Fe Institute Studies in the Sciences of Complexity. Addison-Wesley: Redwood City CA.

Weisbuch, G./Oprea, M. (1994) Capacity of a model immune network. Bulletin of Mathematical Biology 56:899–921.