BioSystems, 11 (1979) 167—182

BOOTSTRAPPING ON THE ADAPTIVE LANDSCAPE

MICHAEL CONRAD

Computer and Communication Sciences, University of Michigan, Ann Arbor, MI, U.S.A.

Different versions of a gene or of a multigenic system may be essentially equivalent so far as the specific function of the structures which they code for or control is concerned, but very different with respect to their amenability to evolution. The structural features which increase evolutionary amenability are a disadvantage to the organism in terms of energy. Nevertheless, they accumulate in the course of evolution as a consequence of hitchhiking along with the desirable traits whose evolution they make possible. This is the bootstrap principle of evolutionary adaptability. In terms of the adaptive landscape bootstrapping corresponds to populations evolving in such a way that they occupy regions of the landscape which are more amenable to evolutionary hill climbing. The bootstrapping idea has implications for structure-function relations in a number of complex biological information processing systems, including biochemical systems, the immune system, and the brain. Bootstrapping is also discussed in connection with the origin of information processing (the origin of life) and in connection with possible designs for macromolecular computing systems.

1. Introduction

Enzymes are switches which turn reactions on and off. One might therefore be tempted to think that all essential aspects of information processing in a chemical medium can be understood in analogy to information processing in an electronic switching network. This is not the case. A subtle, but critical, property of enzymes is that their structure has a dual description which is non-arbitrary. One description is in terms of primary structure, or sequence of amino acids. This description is non-arbitrary because information accumulated in the course of evolution is stored in the form of the sequence of bases in the gene which codes for this primary structure. A second description is in terms of tertiary physical properties, principally nuclear configuration, patterns of electron distribution, and dynamical motions of the entire system. This tertiary description is the description which is natural for understanding the function of the enzyme, what substrates it recognizes, what it does to these substrates, how it responds to control molecules, what other macromolecules it self-assembles with to form more complex structures. This is usually summed up in the idea, to some extent metaphorical, that the function of a protein is determined by its shape. The shape (or tertiary structure) is a second representation of the information accumulated in the primary structure.

This dual representation of information in enzymes is of critical importance for the following reason. The tertiary representation arises from the primary representation through an energy-dependent folding process. Thus by changing the primary structure slightly one can often change the shape and therefore pattern recognition capabilities of the enzyme slightly. The enzyme is thus more than a switch. It is also a specification of what is to be switched. Furthermore, it has the property that this specification can be changed gradually. In this paper I want to concentrate on the significance of this point. The first objective is to show that the gradualism property is the key to the amenability of biological systems to evolutionary optimization. Furthermore, the degree of gradualism can itself be increased in the course of evolution for more effective evolutionary optimization. In a sense

to be made precise it pulls itself up by its own bootstraps. The second objective is to illustrate the role of shape-dependent molecular pattern recognition and the bootstrapping of amenability to evolution in a number of complex biological systems, including biochemical systems, the immune system, and the brain.

Each of these systems is an information processing system. I want to clarify what I mean by this. As I use the term, information processing means selective dissipation of energy. Assuming no advance knowledge of the internal structure of a system, it must be admitted that there are virtually an infinite number of possible pathways of dissipation. Some of these alternative pathways are more interesting or useful than others. Interesting is of course a value judgement. Yet in the broad context of evolutionary theory it can be given a meaning, at least in principle. The pathway has some value if it contributes to reproductive success. Our interest is in the evolution of this small, *selected* set of pathways. In living systems catalysts — principally the protein enzymes — select the pathways and the process of selective evolution selects the catalysts. When one programs a computer one also selects particular pathways of dissipation. The computer does not process information through the action of enzyme catalysts. The living system which programs it does. The thesis of this paper is that this difference is all important for the evolutionary optimization of information processing.

2. Optimization of information processing

2.1. Bootstrap principle of evolutionary adaptability

The unique power of shape dependence is that it allows amenability to evolution to evolve. The effectiveness of evolution by variation and natural selection depends not only on the strategy of search (the heuristics of variation, so to speak), but also on the likelihood that variation produces useful structural and functional changes in the organism. Unlike the search strategies themselves, this amenability depends on the organization of the system, in particular on the physiochemical nature of the genotype-phenotype relationship. In a number of previous papers (Conrad, 1972a, 1974a, 1977, 1978a) the author has argued that this organization, and therefore the likelihood of functionally acceptable variations, is itself a trait which is variable and subject to selection. In short, biological systems are capable of evolving amenability to evolutionary optimization. I call this the bootstrap principle of evolutionary adaptability because variant lines with greater amenability to evolution will almost always be the lines in which new, useful traits appear. This increase in amenability is maintained by the advantageous traits whose likelihood of appearance it increases, thus setting the stage for a further round of increase and maintenance of increase. Even if this amenability is energetically costly (which it always is), it will still appear and be maintained in the course of evolution as long as the adaptively advantageous traits whose likelihood of appearance it increases give an overall energy advantage. To some extent the notion of riding along is analogous to the notion of hitchhiking which has been used in the context of selection for high recombination rates (Strobeck et al., 1976) and earlier in the context of selection for high mutation rates (Cox and Gibson, 1974). In these cases a gene enhancing recombination or mutation hitchhikes along with a fitness-increasing gene belonging to the same linkage group. In the case of a single protein the structural alteration which increases amenability to evolution is on the same gene. A more critical difference is that the hitchhiking effect is stronger for amenability than for recombination and mutation. This is because change in amenability has a more pronounced effect on the rate of evolution, as will be shown.

The principle is also a bootstrap principle

in the sense of arising from the self-contained interplay of Darwinian evolution and the shape-dependence of the macromolecular processes which underlie it at the cellular level, without any reference to the specific details of the system or of the environment. Evolutionary amenability creates itself through this interplay. This is clearest and most fundamental at the level of the relation between the single gene and the protein for which it codes. The folded shape, hence function, of the protein is determined by weak interactions among its constituent amino acids (e.g., van der Waal's interactions, hydrogen bonds, hydrophobic interactions, sulfhydral bonds, among others). The linear sequence of amino acids is in effect the translated genotype of the protein and the features of its folded shape the phenotypic traits which are relevant from the standpoint of natural selection. Suppose that through a process of variation and selection a gene sequence emerges which codes for a protein organized in such a way that it expresses more of the possible mutations in terms of functionally acceptable variations in structure. This will happen if features of the folded shape critical for function change gradually in response to a larger percentage of possible point mutations. There are three important mechanisms. Amino acids may be added to the chain whose principal function is to increase the number of weak bonds which maintain such critical features. The ramifications of a typical point mutation will then be smeared out over a larger structure or absorbed most prominently in features of the shape irrelevant to specific function. Such extra features will in effect buffer the expression of mutation. The second mechanism involves an increase in the number of amino acids with close structural analogs. More replacements will then be possible which produce a small perturbation in the overall structure. In present day systems the number of amino acid types is of course fixed. However, it is still possible to build the protein out of types with a greater number of close structural analogs. The third mechanism involves the use of organizational formats (as in the immunoglobin molecule, to be discussed) which more effectively utilize the weak bonding among a given number of amino acids. All these mechanisms are based on redundancy, either of weak bonding or features of amino acid shape.

The basis idea is thus that two versions of a protein, A and B, may be for all practical purposes identical so far as specific physiological function is concerned, but may be quite different as regards transformability of this function. The rate and binding parameters of these two proteins will be the same, but the typical degree of change of these parameters with single changes in primary structure is quite different. The importance of this difference for the rate of evolution can be seen from the following simple model (Conrad 1972b, 1978a). Suppose that a population carries protein S_0, that S_m is the genetically closest protein of higher fitness, that M independent genetic changes are required to jump from S_0 to S_m, and that all intermediate protein forms are unfit. Then the average number of generations required for the appearance of S_m is given by

$$\bar{\tau}_{0m}^{(m)} = \frac{20^m}{N_0 p^m (1-p)^{n-m}},$$

where (m) indicates the number of required simultaneous genetic events, N_0 is the initial population size, n is the length of the protein, p is the mutation probability, the factor 20 enters because mutation can lead to any of 20 amino acids, the factor $(1-p)^{n-m}$ enters because the $n - m$ remaining amino acids must not change, and $\bar{\tau}_{0m}^{(m)} < 1$ would mean that more than one S_m is expected to appear after one generation. Alternatively, suppose that of the $m!$ possible ways in which S_0 can change into S_m by m single changes in amino acid sequence, there is at least one for which every protein species in the sequence has at least slightly increased fitness (or at least not significantly decreased fitness). The average required number of generations is

now given by

$$\bar{\tau}^{(1)}_{0m} = \frac{20^m}{N_0 p(1-p)^{n-1}} + (m-1)D, \quad m \geqslant 1,$$

where D (the delay time) is the number of generations which it requires for the population to grow to the same size as the old population. The simplifying (and worst case) assumption is that genetic variation is turned off for this number of generations. Even so the rate advantage for the stepwise over the simultaneous mode of evolution is astonishing. For a protein of length 300, a mutation rate of $p = 10^{-8}$, and a step length of $m = 2$, the advantage is of the order 10^9 (assuming $D = 1000$ and $N_0 = 10^6$). If $m = 3$ the advantage is of the order 10^{18}. Variation of protein length, delay time, or population sizes has no significant impact on the relative advantage, which rapidly grows to astronomical proportions as step length increases. It is this consideration which justifies the statement that the possibility of evolution by variation and natural selection depends on the possibility of stepwise evolutionary change and that any variation in the direction of one step gradualism will always be fixed despite the fact that an increase in the number of amino acids, in the number of types, or the introduction of extra organizational features is always costly in terms of energy. This can happen even if both versions of the protein are capable of undergoing one step evolution, but one is faster by a constant factor because its greater degree of gradualism opens up a greater number of single step pathways. Even under these circumstances most favourable evolutionary changes will occur in the more evolutionarily amenable protein and therefore this amenability will always be carried along by selection for these changes. This will stop happening when the protein has evolved to a point where an increase in gradualism is more likely to give rise to loss of selectively valuable traits rather than in improvement; in short, to valley slipping rather than hill climbing.

The idea extends to higher levels of genetic organization. Two possible versions of an organism may be essentially the same so far as their adaptations to the specific environment are concerned, with one energetically less efficient but evolutionarily more amenable than the other. In so far as the first such organisms to appear are more likely to belong to the more amenable line, the amenability characteristic will be fixed. Even if this eventually becomes undesirable — for example, if conservation of highly evolved traits becomes more important — the general organizational mechanisms which allow amenability will almost certainly persist, though possibly in diminished or suppressed form. In the multigene system these mechanisms do not include folding, since folding is a property of proteins coded by one or at most a few genes. However, there are three important alternative mechanisms. The first is reliance on polygenic inheritance. A trait determined by the additive effect of attenuated enzymes coded by redundant genes is clearly more gradually transformable than a trait determined by a single potent gene. The redundancy of genes plays a role analogous to the redundancy of weak bonding in an enzyme, except that the gradualism is now based on change in enzyme number rather than shape (with both affecting the rates of processes). The second mechanism is regulation of gene expression. This can also be used to effect gradual variation of concentrations. The third mechanism involves alternate orderings of genes on the genome. Many alternate orderings are for all practical purposes completely equivalent from the standpoint of development. This makes possible higher level genetic shuffling processes (such as crossing over and recombination) and is therefore the basis for hierarchies of stepwise variation of genes and blocks of genes. These mechanisms depend ultimately, though indirectly, on shape dependence, either because they involve rate control (in the case of quantitative inheritance and regulation) or shape dependent self-ordering processes (which reduce the required constraint on the order of genes on the genome).

Gradualism is a phenomenological fact for many, though not all, proteins (as determined by comparative studies or by experiment, as with dehydrogenases, cf. Rossman et al., 1975; Wills, 1976). Starting with this fact, it is possible to establish the existence of mechanisms for buffering mutation through an entirely different argument. The technique is to study the effect of primary on tertiary structural uncertainty, utilizing the fact that small tertiary uncertainty is a necessary condition for gradualism and that the relation between primary and tertiary structure is unique (in the sense of Anfinsen's unique folding hypothesis). The argument is presented elsewhere (Conrad, 1978b).

2.2 Bootstrapping on the adaptive landscape

It is useful to picture the bootstrap principle in terms of an adaptive landscape, or plot of fitness against genetic structure. The classical adaptive landscape is a plot of fitness against gene frequency and environmental parameters. Here, however, the interest is in molecular structure and thus the axes should be sufficient to specify the detailed sequence of bases in DNA. The most convenient way to do this is to replace the gene frequency axis by a set of axes for specifying the base sequence, with a distinct axis for each possible base at each possible position. The presence of a base of a particular type at a particular position is represented by a one on the appropriate axis and its absence by a zero. The space is thus filled by a hypercloud of lines whose projections on the fitness and nucleotide axes is between zero and one. If all the points in the hypercloud are connected by straight lines, the envelope of these lines becomes a fitness hypersurface with multidimensional peaks and valleys. This is the adaptive surface (or landscape, cf. Fig. 1).

A particular pathway, either up or between peaks, may be easily traversible in the sense that it can be climbed or crossed through a series of single step changes in base sequence which do not involve any significant decrease in fitness; or it may be non-traversible in the sense that multistep changes in base sequence are necessary to jump crevices in fitness. If a system moving on such an adaptive landscape varies in a direction which increases traversibility, it will begin to climb much faster, thus effectively leading to selection for the increase in traversibility. Furthermore, traversibility can always increase in a gradual way since the various redundancies which make it possible can always be increased in a gradual, step-by-step fashion. Increase in traversibility does not mean that new mountains and valleys are created, only that Darwinian evolution inevitably moves the evolving system to (generally higher dimensional) regions of these mountains which are more easily climbable. These portions of the mountains are not as high in terms of fitness because of the extra structure and therefore extra energy required for gradualism. In effect *Darwinian evolution causes biological systems to use free energy to reshape the fine structure of the adaptive landscape into a more negotiable surface.* In some cases selection for this reshaping may be intense. This would certainly be the case if the population must track moving peaks, either because the environment is changing or because other populations are relocating themselves on the landscape.

Selection for amenability to evolution should not be confused with secondary selection involving adjustment of evolutionary search strategies, for example, selection for optimum mutation rate. The former involves the fundamental nature of the genotype-phenotype relation and determines the structure of the space to be searched. The latter involves the algorithm (or heuristics) for searching this space. Improvement of the search algorithm may have a significant quantitative effect on the rate of evolution, but it could never have the dramatic qualitative effect that modifying the structure of the space does.

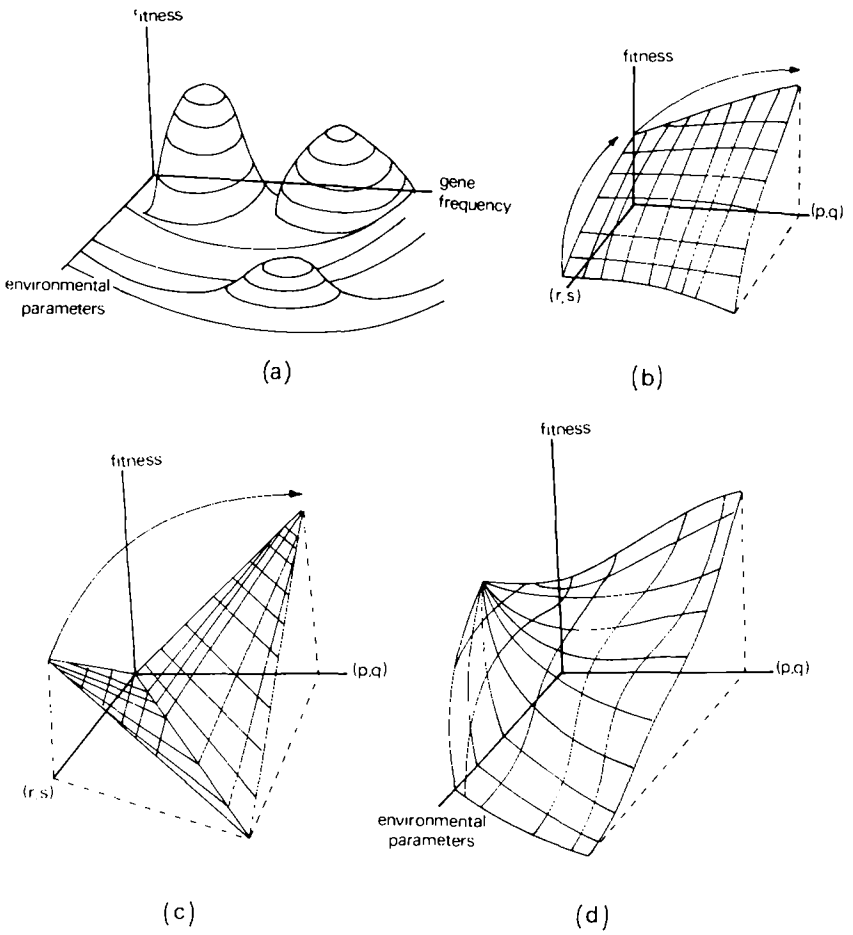

Fig. 1. Classical adaptive landscape and projections of molecular adaptive landscape. (a) Illustrates an environmentally parameterized adaptive surface. (b) Illustrates the envelope of a projection of a molecular landscape on two nucleotide axes (one for nucleotide of type p at position q and the other for nucleotide of type r at position s). The envelope can only assume values of one or zero on these axes, corresponding to the presence or absence of the particular nucleotide at the particular position. The most fit situation is one in which p is present at q and r is absent at s, and the least fit the one where p is present at q and r is present at s. The arrows indicate that it is possible to go from (r,s) to the highest fitness point in two single steps. Note that for a single gene coding for a protein with 300 amino acids 3600 nucleotide axes would be required. (c) Illustrates a two axis projection in which it is not possible, in the projection, to go from (r,s) to (p,q) in a series of single steps since the intermediate steps have zero fitness. A two step jump is necessary. (d) Illustrates the projection of an environmentally parameterized landscape along a single nucleotide axis.

2.3. Significance for optimization of information processing

As is all too well known, many significant optimization problems are extremely difficult to solve and indeed inamenable to any known method of solution, whether deterministic or stochastic. Such problems are of course intractible from the standpoint of any information processing system, whether molecular or conventional. Suppose, however, that it were possible to build an information processing system with the property that it provides significant solutions for optimization problems which are not difficult. The advantage would not be that the system embodied

algorithms more powerful than those now known, but rather that it embodied a greater number of reachable solutions with desirable properties. This is the key to the significance of protein shape for optimization of information processing. What changes during the bootstrapping of evolutionary adaptability is not the search algorithm, but rather the structure of the adaptive surface. By modifying the organization of the protein or the multigene system in such a way that accessibility of optima representing functionally useful patterns of energy dissipation is increased, the bootstrapping process increases the effectiveness of search without modifying the strategy of search.

An argument can be made that this kind of bootstrapping process is not possible in technical information processing systems such as present day computers. Such systems are always constructed from standard types of switching components. Unlike the proteins of biological systems, these standard components cannot be gradually transformed, or transformed at all. The only way such a system could be made more amenable to evolution is by organizing the interconnections among the components differently. The problem is that systems of this type have a powerful feature which is incompatible with a significant increase in evolutionary adaptability. It is possible to effectively communicate to them the rule which generates their behaviour. Thus, such systems are *structurally programmable* in the sense that it is possible to find an algorithmic process which can be used to hardwire any algorithmic process into them. In general, however, this means that it is not possible to predict in advance how different the generating rule will be if the structure of the system is modified. If this were possible, it would be possible to answer the question: will the construction algorithm applied to the generation algorithm (i.e., the algorithm to be hardwired) ever produce a particular hardwired structure? If this question could be answered, it would be possible to determine whether a machine embodying the construction rule and presented with the generation rule will ever halt. However, this would be equivalent to solving the famous halting problem for Turing machines, which is unsolvable.

Any real system is of course finite and the halting problem is only unsolvable in the infinite case. However, what is unsolvable with no finite restriction can in general be expected to be extremely difficult even when a restriction is imposed. In fact, the common experience is that programs (whether or not hardwired) are very sensitive to slight alterations. The effect of such alterations are either unpredictable (otherwise the halting problem would be solvable) or they have no effect at all (because of fault tolerant design). In general it is impossible to define a concept of gradual transformability for systems built on the standard parts principle, hence impossible to bootstrap evolutionary adaptability.

There exists an apparent counterexample to the above argument which is instructive. Imagine a structurally programmable system which consists of a large number of subsystems, many of which are capable of doing something useful, and all of which are connected to each other by at least one sequence of systems with the property that each member of the sequence performs a function only slightly different than its neighbors. Variation and selection, involving activation and de-activation of subsystems, would be an extremely efficient search procedure in this case. The example clearly begs the question of the origin of all the necessary subsystems and is unreasonably inefficient since it requires all of them to be present. If the system were built out of macromolecules instead of standard building blocks, both of these problems would disappear since each member in the sequence could be derivable from its neighbor by a small structural change. Furthermore, each particular system is potentially more efficient in this case since the components can be adapted to the specific task.

The difference between information processing in macromolecular and technical

information processing systems can be thought of in terms of a trade-off between amenability to evolution and programmability. By sacrificing structural programmability, one gains the possibility of efficiently generating a family of special purpose information processing systems by an evolution process. I want to emphasize that this trade-off should not be taken to imply that evolutionary search methods are inefficient for the computer solution of optimization problems. For appropriate adaptive surfaces such methods are often most efficient (cf. Bremermann, 1962; Holland, 1975). Rather it implies that the adaptive surface of a structurally programmable system is itself not suitable for evolutionary search. Only structurally non-programmable systems (such as macromolecular systems) have the property that evolutionary search transforms them in such a way that their adaptive surface becomes more suitable for evolutionary search. This is the subtle key to the success of evolutionary search in producing so remarkable a variety of functional life forms.

3. Role of bootstrapping in complex systems

3.1. Biochemical systems

As a first example of the significance of bootstrapping we return to the initial discussion of information processing as selective dissipation. Protein enzymes select the pathways of dissipation and the process of evolution by variation and natural selection selects the proteins. In the bootstrapping process the expression of point mutation and other genetic variation on protein structure and function is buffered in such a way that function change is more gradual. This means that the rate constants and other parameters (binding constants, mobilities, cooperativities) which characterize the enzyme are more likely to change more gradually in response to point mutation. Picturing a biochemical system as described by a set of rate equations, genetic changes will often give rise only to small perturbations of these equations in the form of changes in the rate parameters. To the extent that the system is structurally stable to these changes, its dynamical behaviour will only change quantitatively, not qualitatively. The structural stability of the folding process is in effect inherited at the level of biochemical dynamics. The bootstrapping of a mutation buffer system in a single protein thus leads to a global bootstrapping of gradualism. Such global quantitative variation is the dynamical analog of the continuous mappings between related species discussed by D'Arcy Thompson (1917). A discussion in terms of similarity of dynamical systems has recently been given by Rosen (1978). At this more global level one is clearly dealing with multigene systems. Structural stability to forms of genetic variation other than mutation becomes important, thus allowing for multigenic bootstrapping (Bargiello and Grossfield discuss such stabilizing mechanisms in this issue). The folding of the protein is of course in many cases teratologic. Even if it is not, slight change in a rate constant could still lead to a "catastrophic" change in the dynamics of the system) (cf. Thom, 1970). This is certainly possible and such radical changes might play an important role in some instances (as has sometimes been suggested for the origin of phyla). However, as in the evolution of a single protein, the vast majority of radical changes in the dynamics of a complex system are unlikely to be viable. In both cases, the evidence is that evolution proceeds by small steps, with gradual transformation of structure and function and through the gradual assumption of new functions by given structures (i.e., transfer of function).

The bootstrapping of evolutionary adaptability has an important implication for the controversy over the relative importance of natural selection versus neutral genetic variation as a determining factor in the course of evolution. A gene or multigene system organized in such a way that the functional expression of its variation is buffered will

inevitably exhibit variations which are neutral or practically neutral as regards selection. However, this kind of organization is a necessary condition for natural selection to be an effective agent in organic evolution. The consequence, only seemingly paradoxical, is that those phenomena which are ordinarily taken as evidence for a neutralist view are really indications that the system has bootstrapped itself to an organization which fulfills conditions necessary for selection to be effective.

3.2. Origin of life

The bootstrapping idea applies in an intriguing way to prebiological evolution. Suppose enzyme A catalyzes the production of enzyme B, which catalyzes the production of C, and so forth. Eventually some enzyme in the series may catalyze the production of A, thus giving rise to a cycle. If such a cyclic process appears it will persist unless it is replaced by another cycle, possibly emerging from a mutation affecting one of the proteins. If each of the enzymes is replaced by an autocatalytic process the system would correspond to a hypercycle of the type proposed by Eigen (1971) as a model for prebiological evolution. Any variant cycle which reproduces more effectively than the original cycle should capture more of the monomers and should thus begin to dominate the pathways of dissipation. Here it is convenient to think of any such cycle as a series of primary structures. One can then think of each enzyme as a primary sequence which codes for the primary sequence of the enzyme whose production it catalyzes (where enzymes can here be taken as including both proteins and nucleic acids). At first there may exist as many codes in such a system as enzymes; the code is entirely non-universal. Cycles in which there is some specialization (such as nucleic acid specialization for information storage) should be more efficient and thus systems will develop in which there are fewer coding rules and eventually only a single coding rule.

The assumption on which this scenario depends is that between the original cycle and one which embodies a single coding rule there is at least 1 sequence of cycles which is easily traversible. Easily traversible means that only one or at most two mutational events separate each cycle in the sequence. This will not be likely if the enzymes in the cycle are significantly changed by mutation. Under these circumstances, most new cycles will be less efficient than the original and eventually the system will be trapped in an undesirable valley. However, those cycles which vary in the direction of greater gradualism will almost inevitably be the ones from which more efficient cycles derive. Thus, in general cycles will bootstrap in such a way that enzymes develop the buffering property. The development of this property is equivalent to the movement of cycles to regions of the adaptive landscape which are more amenable to hill climbing by cycle budding (the cycle or hypercycle mechanism of variation). Once in this region the possibility of finding an easily traversible pathway to a cycle with a coding system in the modern sense is increased. The code need not be the modern code. Once a simple coding system appears it can evolve by a multiple coding mechanism in coenocytic type organizations, or possibly by other mechanisms. In the multiple coding mechanism (Conrad, 1970) a system with a variant code which may initially lead to the production of mostly useless proteins is maintained through a symbiotic relationship with a precursor coding system if it produces at least one useful protein. If such a variant code is superior it will eventually become the dominant coding mechanism.

3.3. Immune system

The most extreme and clearcut example of a system for buffering genetic variation at the macromolecular level is provided by the immune system. Figure 2 illustrates the well-known lobster-like structure of the immunoglobin molecule. The hinges on the arms

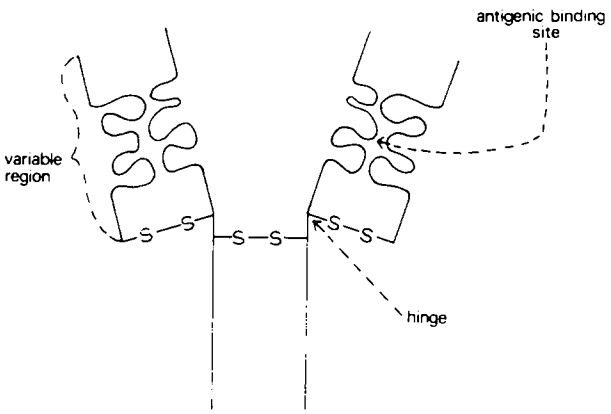

Fig. 2. Schematic illustration of loopings in the "lobster's claws" of the immunoglobin molecule (cf. Roitt, 1974).

allow for a choice of different possible orientations of the antigen binding sites. The feature to note is the loopings of the amino acid chain within the claws. It is the variations of amino acid sequence on these claws (in the so-called variable and hypervariable regions) which allow for the enormous range of antigenic specificities. For any particular antigenic specificity one could imagine a less expensive antibody without loopings. However, such a molecule would be less gradually transformable and would be capable of supporting a smaller range of possible antigenic specificities. The pathways connecting functionally useful specificities would consequently be much more difficult to traverse. The loopings are a grossly observable example of a redundancy dependent buffering mechanism in a protein. The hinged claws are an example of a global format which allows for the most effective use of such redundancy-based buffering mechanisms.

The loopings allow for effective optimization of antigenic specificity. The optimization may be on a phylogenetic time scale, with natural selection acting on inherited variations of the antibody gene. Or it may be on an ontogenetic time scale, as in the well known clonal selection mechanism. In this case the variation is presumably based on recombining inherited parts (germ line mechanism) or on somatic mechanisms of variation, or most probably on some combination of these (cf. Bremermann, this issue).

3.4. Brain models

The pattern recognition capabilities of the immune system are based on the shape-dependent interactions of antibody and antigen. The optimization of these capabilities depends on looping and other buffering mechanisms which give the antibodies great evolutionary adaptability, in both the ontogenetic and phylogenetic sense. An analogy to the brain has often been observed. Optimization of antigenic pattern recognition is a form of learning which, at the systemic level, allows organisms to solve problems associated with their integrity as individuals and with violations of this integrity. Storage of these solutions, hence memory, is essential. Allocation of storage space, time for retrieval of stored memories, and control over their dynamical expression are critically important.

The analogy to the brain is clearly limited. The principal problems the two systems are required to solve are different. The patterns to be recognized by the brain are encoded in patterns of receptor activity or in patterns of neural firing, whereas the patterns recognized by the immune system directly involve the molecular geometry of the antigen. The output must also be expressed in terms of patterns of neural activity and ultimately in terms of behavioral acts of the organism. The associated problems of storage allocation, retrieval and control are present, but necessarily differ in important respects. Nevertheless it is possible to construct models of the brain in which neurons and their macromolecular constituents are so organized that the number and accessibility of useful solutions is increased in fundamentally the same sense that it is increased in the immune system.

The basic requirement of any such model is that the specificity of macromolecules, a surface property, must be coupled to the

recognition of specific patterns of input from other neurons or from receptors. This is possible if different locations on the neuron membrane are potentially activated by different patterns of dendritic inputs, but in fact are activated only if a macromolecule which makes it possible for activation to initiate impulse formation is located there. Such a macromolecule (to be called an excitase) is at present hypothetical. However, the assumption of its existence leads to models in which networks of neurons are amenable to the evolutionary optimization of information processing in the same highly effective fashion that the gene, genetic systems, and the immune system are. Furthermore, it cross-correlates in a natural way with molecular structures (gating molecules) whose existence is supported by experimental evidence and whose specificity is believed to be responsible for initiation and propagation of the nerve impulse (Hille, 1976). The presence or absence of excitases would then correspond to the presence or absence of gating-current sensitive molecular gates (or essential components of such gates) which control the transport of ions across the membrane. Excitases with different primary structures differ with respect to which location they can recognize and bind to. Such locations may be distributed over the cell surface of microneurons with geometrically very distributed inputs, concentrated in conventionally excitable areas (the axon hillock), in individual dendrites, or in dendrites connected by short circuits. The initiated impulse may lead to cell firing or may amplify a pattern of inputs which eventually leads to cell firing. The neuron itself is conceived as a system whose morphological and conductive asymmetries serve to break the symmetry of the dendritic inputs, allowing different patterns of input to lead to different levels of excitation in different locations. The neuron will thus fire in response to the set of input patterns which appropriately excite locations to which excitases are bound. This set can be changed in a step by step fashion by adding or deleting different types of excitases to the system.

The structure of the network stays the same. The morphology of the neuron stays the same. But the input-output characteristics of the neuron are allowed to change gradually. The mechanism of gradualism is either based on the buffering of folded shape (as for single genes generally) or on independence of gene action (one mechanism of multigene gradualism). If the folding of the excitases is buffered by redundancies, such as those which buffer the folding of the immunoglobins, slight changes in primary structure will often lead to excitases which locate themselves in a slightly different region of the membrane or which respond slightly differently to different degrees of excitation. If the excitases did not act independently, the set of patterns recognized by the neuron would in general change radically with the addition or deletion of an excitase. Thus the independence property allows for step-by-step selection of the basic patterns to be recognized and the buffered folding allows for fine adjustment of the ability to recognize each of these patterns. Note that the morphology and conductivity properties of the neuron could also be changed gradually, but in this case any excitase whose function it is desirable to preserve would also have to be modified. The advantage of an excitase based system is that it allows for both gradualism and for highly specific responses to patterns of input.

Two types of optimization processes are possible. The first is evolution by variation and natural selection. In this case the excitases are coded by genes which are inherited from the parent. The second possibility is an ontogenetic learning process analogous to evolution. Two specific mechanisms are possible. The variation may involve mechanisms analogous to those used in the immune system, i.e., may be based on some combination of germ line and somatic mutation mechanisms. If the excitases coded for by such variant excitase nucleic acid is useful, this nucleic acid would be transferred to corresponding neurons in neighboring networks of equivalent structure. Because of the equivalence of structure, the excitases have the same function

in the neurons which they transform. Usefulness is actually not a property of individual excitases, but rather of the network. The new excitase is useful if the local network which contains it leads to more satisfying interactions with the environment when the organism's behavior is controlled by this local network rather than by other networks of the same structural type. The second possible mechanism involves inductive processes analogous to those which play a role in development. The variation would be based on the random production of substances which induce the appearance of an excitase. If this excitase is useful (in the sense defined above) the inducing substance would be exported to corresponding neurons in neighboring networks, where it would induce the appropriate excitases. The two mechanisms (immune system analogy and developmental analogy) are isomorphic as optimization processes. However, the developmental mechanism requires a great deal of pre-existing correlation between inducers and excitases and also does not allow for as large a repertoire of excitases.

The model outlined above (called the evolutionary selection circuits model) has been described in detail elsewhere (Conrad, 1974b,c). Such models allow for efficient evolutionary optimization of pattern recognition and generation. They are efficient in terms of number of neurons required for a task since each neuron can be adapted specifically for a desired task. The selection circuits model deals with gradual learning, not with memory acquisition, manipulation, and retrieval. Any device with general powers of information processing must have these capabilities. The model can be interfaced with another model, called the reference neuron scheme, which does incorporate these capabilities. The memory manipulation capability is based on conformation changes in dendrite-bound receptor molecules. The reference neuron scheme and its interfacing with the selection circuits model is described elsewhere (Conrad, 1974c; 1976a,b). It should be noted that other authors, in particular Smith (1962) and Hyden (1967), have presented alternative discussions of aspects of brain function involving inductive mechanisms. More recently, Edelman (1978) has presented a model of brain function which is selectionist in the evolutionary and immunological sense, though not based on the action of individual macromolecules.

3.5. Molecular computers

Someday it may be possible to construct artificial information processing systems which utilize the essential principles that make evolutionary optimization possible in living systems. Such a system would require numerous replicas of networks of modules, each of which can serve as a symmetry-breaking neuron. Inputs can come from, say, photosensitive receptors, while outputs could be coupled to a process to be controlled or might be used for classifying input patterns. If excitases control a reaction which initiates the transmission of signals from module to module, it would be possible to build devices which would be capable of efficiently recognizing patterns of input and generating appropriate patterns of output in response. One would not know in advance what primary structures the excitases should have in order to migrate and bind to appropriate sites. However, many excitases could be generated using genetic engineering techniques and then embedded in the network, with the appropriate choice being determined by an evolution process. Once a network with the desired properties is produced, many copies could be fabricated, using the same preparation procedure. The system would be adaptive on its own if the mechanisms of evolution were included in its operation.

The fabrication of such a system involves many technical problems which are difficult and which I do not want to discuss here. The critical point is that such an enzymatically controlled network allows for a rich set of reachable functional realizations. Each real-

ization is a special purpose device and the evolution process makes it possible to find any of a large number of potentially useful devices. As an analogy, one can think of a telephone network, with banks of telephone booths and lines connecting all the telephones. The lines correspond to the connections between cells, the banks of booths correspond to the cells and each booth to a location on the membrane. Changes in the structure of the system — in the connectivity or architecture of the booths — are really not too essential to the processing of information. What is really essential are the decisions which humans using the phones make. In our molecular computer the excitases replace the humans. They are using all the complex structures of the cell and the connectivity between cells to communicate with one another. They are clearly much more elementary decision makers than humans are. But then so are the macromolecules responsible for the basic decision processes in the human brain.

4. Conclusions

All the macromolecular information processing systems discussed in this paper — prebiological cycles and hypercycles, multigenic biochemical systems, the immune system, the brain, designs for macromolecular computing systems — share a single, basic feature. The addition of mechanistically superfluous structure allows for increased evolutionary adaptability. By making the structure more expensive, for example by adding to the redundancy of weak bonding in an enzyme or by increasing the number of genes necessary to produce a structure, the degree to which function changes with genetic change can be made more gradual. Such addition of mechanistic superfluity does not contradict natural selection. It hitchhikes along with the desirable traits which it makes possible. Energy is required to maintain these superfluities. This energy in effect makes the adaptive landscape more amenable to evolutionary hill climbing. Another consequence is that the relation between structure and function becomes more difficult to understand. In part this is because the structures are more complicated. More important, it is because the relation between structure and function must be such that a small perturbation of the structure often gives rise a small perturbation of the function. This is not in general the case if the switching elements in the system are standard components which cannot themselves be gradually modified. While systems made from gradually modifiable components cannot be as efficient as systems made of appropriate unmodifiable components, they are in general more efficient than systems made of arbitrarily preselected unmodifiable components. This is because unmodifiable components are in general not specifically adapted to the required task. In biological systems the components specifically adapt to the task, but something must be paid for this ability to adapt.

These considerations have an implication for genetic programming. In principle there should exist some DNA sequence which would code for an organism with all the capabilities of a human being, yet which would be much more efficient. It would be more efficient if the evolutionary adaptability of individual genes and of the genome as a whole were forfeited. But such a genome would correspond to an adaptive peak which is so isolated that it could never be discovered by any optimization process, including optimization by variation and natural selection.

Acknowledgement

I would like to thank J. Grossfield and T. Bargiello for extremely helpful comments on the manuscript.

References

Bremermann, H.J., 1962, Optimization through evolu-

tion and recombination, in: Self-organizing systems, Yovits, Jacobi and Goldstein (eds.) (Spartan Books, Washington, D.C.).
Conrad, M., 1970, A mechanism for the evolution of the genetic code, Currents in Modern Biology (now BioSystems) 3, 260.
Conrad, M., 1972a, Information processing in molecular systems. Curr. Mod. Biol. (now BioSystems) 5, 1.
Conrad, M., 1972b, The importance of molecular hierarchy in information processing, in: Towards a theoretical biology, Vol. 4, C.H. Waddington (ed.) (Edinburgh University Press, Edinburgh).
Conrad, M., 1974a, The limits of biological simulation. J. Theor. Biol. 45, 585.
Conrad, M., 1974b, Evolutionary learning circuits. J. Theor. Biol. 46, 167.
Conrad, M., 1974c, Molecular information processing in the central nervous system, parts I and II, in: Physics and mathematics of the nervous system, M. Conrad, W. Guttinger and M. Dal Cins (eds.) (Springer, Heidelberg).
Conrad, M., 1976a, Molecular information structures in the brain. J. Neurosci. Res. 2, 233.
Conrad, M., 1976b, Complementary molecular models of learning and memory, BioSystems 8, 119.
Conrad, M., 1977, Evolutionary adaptability of biological macromolecules. J. Mol. Evol. 10, 87.
Conrad, M., 1978a, Evolution of the adaptive landscape, in: Theoretical approaches to complex systems, R. Heim and G. Palm (eds.), Springer Lecture Notes in Biomathematics No. 21 (Springer, Heidelberg).
Conrad, M., 1978b, Mutation-absorption model of the enzyme. Bull. Math. Biol. in press.
Cox, E.C. and T.C. Gibson, 1974, Selection for high mutation rates in chemostats. Genetics 77, 169.
Edelman, G.M., 1978, Group selection and phasic reentrant signaling. A theory of higher brain function, in: The mindful brain (The MIT Press, Cambridge, Mass. and London, England).
Eigen, M., 1971, Self organization of matter and the evolution of biological macromolecules. Naturwissenschaften 10, 465.
Hille, B., 1976, Gating in sodium channels in nerve. Ann. Rev. Physiol. 38, 139—152.
Holland, J., 1975, Adaptation in natural and artificial systems (University of Michigan Press, Ann Arbor, Mich.).
Hyden, H., 1967, Biochemical changes accompanying learning, in: The neurosciences, G.C. Quarton, T. Melnechuk and F.O. Schmitt (eds.) (Rockefeller University Press, New York).
Roitt, I., 1974, Essential immunology, 2nd edn. (Blackwell Scientific Publications, Oxford).
Rossman, M.G., A. Liljas, C. Branden and L.J. Banaszak, 1975, Evolutionary and structural relationships among dehydrogenases, in: The enzymes, 3rd edn., Vol. 11, P.D. Boyer (ed.) (Academic Press, New York).
Rosen, R., 1978, Dynamical similarity and the theory of biological transformations. Bull. Math. Biol. 40, 549.
Smith, C.E., 1962, Is memory a matter of enzyme induction? Science 138, 889.
Strobeck, C., J. Maynard Smith and B. Charlesworth, 1976, The effects of hitchhiking on a gene for recombination. Genetics 82, 547.
Thom, R., 1970, Topological models in biology, in: Towards a theoretical biology, Vol. 1, C.H. Waddington (ed.) (Edinburgh University Press, Edinburgh).
Thompson, D'Arcy Wentworth, 1917, On growth and form (Cambridge University Press, Cambridge England).
Wills, C., 1976, Production of yeast alcohol dehydrogenase isoenzymes by selection. Nature 261, 26.

Commentary

T. Bargiello and J. Grossfield

This paper presents a convincing argument which establishes the significance of gradualism to evolutionary processes and the mechanisms by which gradualism can evolve. The implications of the theoretical considerations developed in this paper have a possible bearing on a number of diverse and seemingly unrelated areas.

The existence of gradualism depends in part, on the ability of biological systems to produce "mechanistically superfluous structures". At the protein level, this may be accomplished by the three mechanisms described by Dr. Conrad. The properties of these variant proteins are interesting in that they are most likely less efficient energetically and hence most probably reduce individual fitness. However, the selective advantage of amenable proteins lies in their ability to generate future mutants which have properties similar to their own. This latter aspect results in gradualism and ensures the potential for rapid increases in population fitness. The fitness considerations described above occur over two different time courses. The mutations required to produce an amenable protein

structure are selected in the short term, whereas the subsequent evolutionary advantage of proteins with an amenable structure is clearly a long term process. The limiting factor to the evolution of gradualism appears to be the likelihood of the occurrence and maintanance of the first mutation which results in the production of an evolutionarily amenable protein structure.

On the assumption that such a mutation will eventually occur during the evolutionary history of the population, it is evident that proteins characterized by a structure amenable to further evolutionary change will become prevalent in natural populations. This forms the basis of a convincing argument favoring the existence of selectively neutral genes in populations. In this sense, the evolution of protein structure can be viewed as an adaptation to tolerate the perturbations of future mutations. Therefore, if changes in structure are translated into small functional differences (effectively neutral) then one might predict that a correlation exists between the amenability of a protein to further evolution and the number of alleles at that locus. Gene loci characterized by a large number of alleles would be expected to have evolved an amenable structure, whereas those which are monomorphic (unamenable) would give rise to alleles with such radically different catalytic properties that they would be rapidly eliminated from the population. These considerations could conceivably be extended to explain the apparent correlation between protein structure and genetic variability (cf. Zouros, Nature 262, (1976) 227 and Koehn and Eanes, Theor. Popul. Biol. 11, (1977) 330). The existence of such a correlation depends partly on the question of whether all proteins necessarily have a functional structure which can be made amenable to further evolutionary change. If the answer to this question is yes, then it is apparent that all biochemical systems will eventually become amenable, and the possibility of neutral mutations will be determined by the stability of the translation of the primary protein structure to its tertiary structure with respect to amino acid substitutions. The frequency of neutral mutations at a particular genetic locus coding a protein would be determined by the evolutionary age of the protein together with the likelihood of the existence of an amenable structure.

Conrad's response to Bargiello and Grossfield

Bargiello and Grossfield correctly point out that the limiting factor in the evolution of gradualism is the likelihood that an amenability-conferring mutation will occur and be maintained long enough to allow for the appearance of a fitness-increasing trait. Such a double event might appear unlikely. However, a gradualism-conferring mutation, for example the insertion of one or a few extra amino acids into the protein, may be only modestly costly in terms of energy. A certain fraction of such evolutionarily enhanced proteins will always be present. Bargiello's and Grossfield's statement about the limiting factor can thus also be thought of in terms of the chance of a fitness-increasing mutation appearing in the smaller, evolutionarily-enhanced fraction of the population as compared to its chance of appearing in the larger, unenhanced fraction. If the process is repeated frequently enough — where frequently means there is another round before the increased amenability falls off — the degree of gradualism would continue to build despite the fact that the load in terms of individual fitness might eventually become significant.

F.H. Kirkpatrick

Dr. Conrad presents an illuminating argument for an evolutionary selection mechanism for proteins whose tertiary structure is slightly "unstable" to small variations in amino acid composition. Hemoglobin provides an elegant example of such a protein, and cytochrome *c* to a certain extent is an example of the postulated end-state: a protein so far evolved that no major functional variations are encountered, with function almost independent of about 2/3 of the residues and absolutely dependent on the rest.

This model helps explain one of the more baffling phenomena in current work on intracellular regulation: the ubiquity and multiplicity of signal systems involving cAMP, cGMP, calcium and protein phosphorylation. At the moment, it is possible to explain any one of a variety of forms of regulation by one form or another of these general mechanisms; but it is difficult to envisage how all of them can be operating simultaneously. A cell will typically have dozens of phosphorylated proteins, when carefully analysed, and most of these reactions appear to be regulated via cAMP, cGMP and/or intracellular calcium. It is not clear why cyclic pyrimidines and cyclic deoxynucleotides have not been employed to provide more precise control.

The "bootstrapping" principle helps explain why the existing modes of regulation have been selected. Phosphorylation has been selected as a *reversible* post-translation modifier of conformation, which can be applied to a variety of proteins without requiring modification of their primary sequence. Instead, the sequence of the protein kinase (or of its regulators) can be varied. It seems clear that a cell with this capability is more amenable

to bootstrapping than one without it. Such a cell will rapidly develop a variety of protein kinases and phosphoprotein phosphatases with variable sensitivity to regulation. Since some of the kinases will be "sloppy", the bootstrapping cell will have to develop ways to *prevent* phosphorylation of critical sites. (This could explain the existence of methylhistidine in skeletal actin and myosin, but not in many cytoplasmic actins and myosins, since histidine is a phosphorylatable residue (M. Weller, Biochim. Biophys. Acta 509, (1978) 491—98).

Such a cell is now a system in which control is more "analog" than "digital": the average level of phosphorylation of each subsystem is now the critical variable, rather than its presence or absence, and this in turn can be finely regulated by only a few input signals provided that the regulator is responsive to most or all of those inputs. Such a regulatory system can be adjusted in very small increments, i.e., should bootstrap well. Viewed in this light, systems using larger numbers of regulators may have been crowded out because they would have had to spend more energy in setting up such a system without gaining any advantage in ability to bootstrap.

The dichotomy between structurally programmable and structurally non-programmable systems is useful for thinking about structures of systems, but I think it does not apply neatly to the real world. The consequences of changing one amino acid in a protein are, in general, no more predictable than the results of changing one character in a computer system. Like computers, biological systems are at each level composed of standard building blocks (atoms, amino acids, nucleotides, etc.) which are not especially modifiable. The principal problem in building computers capable of evolution is that their builders rigorously select for absolute programmability. If there were an economic force favoring systems that were only 90% predictable, could not evolution of the biological sort occur?

Conrad's response to Kirkpatrick

Dr. Kirkpatrick has provided a valuable and (to me) unexpected example of a system which has just the kind of organization one would expect on the basis of bootstrapping. His discussion suggests a vantage point for appreciating other regulatory systems. The example makes very concrete the perverse point that only systems which are hard for the biologist to analyze are likely to predominate in the course of evolution.

Dr. Kirkpatrick also raises an interesting point about the distinction between structurally programmable and structurally non-programmable systems in the real world. I would argue that the distinction depends on whether it is more natural to think of a system as containing a finite or infinite number of possible building blocks. Atoms and amino acids are building blocks, but the natural analogs of the switching elements in a computer are proteins and other macromolecules. There are a finite number of possible proteins, but the number is so large that it seems to me to be more useful to think of biological systems in terms of infinite rather than finite programming. My argument is that it is only sensible to put a metric on the degree of function change in the former situation, not that the detailed change could be predicted. In other words, the detailed consequences of an amino acid change are in general unpredictable, as Dr. Kirkpatrick points out, but nevertheless it makes sense to talk about the idea of buffering the effects of this change.