Evolutionary Systems, Gertrudis van der Vijver, Stanley N. Salthe
and Manuela Delpos (eds) (Dordrecht, Kluwer, 1998): 129-140.

Information Increase in Biological Systems: How does Adaptation Fit?

John Collier
pljdc@alinga.newcastle.edu.au
7 June 1996

Abstract

Progress has become a suspect concept in evolutionary biology, not the least because the core concepts of neo-Darwinism do not support the idea that evolution is progressive. There have been a number of attempts to account for directionality in evolution through additions to the core hypotheses of neo-Darwinism, but they do not establish progressiveness, and they are somewhat of an ad hoc collection. The standard account of fitness and adaptation can be rephrased in terms of information theory. From this, an information of adaptation can be defined in terms of a fitness function. The information of adaptation is a measure of the mutual information between biota and their environment. If the actual state of adaptation lags behind the state of optimal adaptation, then it is possible for the information of adaptation to increase indefinitely. Since adaptations are functional, this suggests the possibility of progressive evolution in the sense of increasing adaptation.
Keywords: evolution, information, adaptation, fitness, entropy, progress

1. Introduction

The notion of progress has been virtually banned from mainstream contemporary biology (Nitecki, 1988: viii), even though it still lurks in the background (Ruse, 1988). This is a marked change from earlier views, which saw evolution to be distinctly progressive. The complete explanation of this change is no doubt complex, involving factors dear to historians and sociologists of science, such as the breakdown of the Enlightenment view of science and society, and the education and class commitments of prominent authorities. It is probably too early to historically evaluate these forces with any sort of objectivity. I will focus instead on the internal logic of the shift, and the extent to which it is justified: Progress has no theoretical role in contemporary neo-Darwinian orthodoxy. As John Maynard Smith (1988) points out, given two states of a biological system, there is nothing in Fisher's "fundamental theorem of natural selection" that would allow a biologist to determine which state is earlier. The hardening of the Modern Synthesis of genetics and population biology has permitted the gradual realisation that directed processes of any kind, let alone progressive ones, are neither a probable consequence of nor presupposition of its core theses alone.

This situation is not entirely satisfying, since it seems evident that there has been an increase in *something* from the earliest life-forms to those we find today (Ayala, 1988), and the neo-Darwinian orthodoxy tells us that traits that have not been eliminated are more adaptive, and hence more *functional*, than those that are. Given that functionality is a teleological notion, and that adaptive traits are the only neo-Darwinian candidates that can support an increase in *anything*, it isn't a great step to assume that biological progress is just the increase in functionality resulting from adaptive processes. Indeed, this has been a popular view among prominent biologists (Ruse, 1988; but see Williams, 1966). This consequence can be avoided by admitting that directional changes *are* found in the evolutionary record, but that these changes do not involve increased overall functionality (nor, perhaps, are they directed by increases in any global property). Furthermore, directional change is arguably not global, but local only (Hull, 1988).

Evolution might be locally, or even globally, directional, but it isn't progressive. There are several ways to support this move without doing (too much) violence to orthodox neo-Darwinism.

On the other hand, the problem of directionality, and other empirical problems with the Synthesis have led to calls for extending, expanding or replacing the neo-Darwinian paradigm (e.g., Eldredge, 1985; Wicken, 1987; Brooks and Wiley, 1988; Wesson, 1993; Salthe, 1993). Behind these calls is a recognition that evolution is an historical process, in which historical contingencies constrain later evolution, leaving their trace in contemporary observable traits. This implies at least an overall *direction* of evolution, if not an increase in some property over time (Ayala, 1988). In this chapter, I will show how the implications of the historical constraints on evolution can account not only for directed adaptation, but also for genuine increases in adaptation, and hence progressive evolution. The basis for this is a more precise notion of biological form that permits us to make comparisons of the "fit" of traits (and of the organisms that bear them) to their environment.

2. Adaptation produces information

Although everything must have some form, morphology and morphogenesis have been especially important in our understanding of biological systems. In the past, our understanding of biological form has been largely qualitative. Now, information theory allows us to precisely quantify (if not measure) form. Perhaps the most salient feature of biological form is its organised complexity. Although individual subsystems might expire, the typical pattern we observe in biology (whether in evolution, ecology or development) is growth of complexity and organisation across an integrated system. In principle, this can be measured as an information increase, either in terms of computational complexity or computational depth, or (more likely) both.

One important aspect of (to speak more definitely, a *component* of) biological information is adaptedness. Adaptedness is the fit of an organism with its environment that increases the survival potential of the organism's genes. Adaptation is the process by which adaptedness is produced and/or maintained through natural selection. An adaptation, in turn, is a trait that is fit, and exists because of adaptation. I will give an account of the information involved in adaptation that is based in orthodox accounts of fitness. This account involves the comparison of the information in a trait with the information in the environment, i.e., between the form of traits and the structure of its environment. I will then discuss how that information can increase. My account is compatible with a physical interpretation of information that makes it measurable, at least in principle, through direct intervention, like any other scientifically useful physical property. My account also makes increase in the information of adaptation a likely if not inevitable outcome of evolutionary processes. We can keep the core of selection theory, while reconceiving it in a way that makes information increase a likely outcome of adaptation. This approach can, at the same time, lay the ground for the integration of adaptation with other biological processes involving information transmission and transformation.

3. Genetic determination of inheritance

Natural selection is environmentally caused differential reproductive success resulting in fitness differences[1]. Selection can act on populations, organisms or genes. I will focus on the

[1] Rather than selection by chance or deliberate intervention. The former is called neutral (Kimura, 1983), while the latter is called artificial selection (but see Brandon, 1990).

fitness of organisms and their traits for convenience, but much of what I say can be applied, *mutatis mutandis*, to populations and genes. The fitness of a trait of an organism is a measure of the trait's adaptive value for the organism, relative to the organism's environment. A trait has adaptive value if it serves the survival of the lineage of which the organism is a member. The complete fitness of an organism is the net effect of the fitnesses of its traits. Since survival is necessary for the continuing existence of lineages, and evolutionarily significant traits are inherited, the existence of some traits can be explained by their contribution to adaptive value, or fitness (Wright, 1973; Sober, 1984). A trait is biologically functional just insofar as it contributes to the fitness of the organism that possesses the trait.

It is doubtful that fitness can be reduced to lower level physical properties (Sober, 1984). This means that fitness must be defined in terms of its effects on survival. On the other hand, it is impossible to give an adequate definition of fitness entirely operationally (Brandon, 1990). In order to avoid circularity, the fitness of a trait should be defined as a propensity for the organism that has it to produce lineages of organisms that also have the trait, rather than being measured in terms of the number of replicas of the trait actually produced (Mills and Beatty, 1979). (Roughly, propensity is an inherent tendency that something will occur with a certain probability. I will have more to say below.)

Traits that are not inherited, either directly or indirectly, cannot be fit, even though they might contribute to the survival and reproduction of the organism that has them. The reason for this is that they can play nothing but an accidental role in evolution. Only inherited traits can be the cause of a consistent trend. The effects of non-inherited traits last for only one generation, and have no regular effect on the long term survival of a lineage. Inheritance is not an all or nothing phenomenon: A trait might have an environmental and a hereditary component. Only the latter matters to patterns in evolution. The main form of inheritance is genetic. A trait is *genetically determined* if it "differs from other traits at least in part as the result of the presence of one or more distinctive genes" (Wilson, 1978: 19). Genetic determination does not preclude other forms of inheritance, which can be cultural, environmental or somatic. The totality of all means of inheritance is called a *developmental system* (Griffiths and Grey, 1994). In what follows I am going to simplify things by considering only the genetically determined part of inheritance. This permits a reduction of the information in traits to the information in their genetic determinants. Information in a developmental system that is stored environmentally or culturally will complicate the story considerably, but I will rely on the assumption that the pattern of explanation will be the same. Neo-Darwinism generally assumes that genetic determination is the only evolutionarily significant form of hereditary determination, so the simplification is not unusual.

4. Information of adpatation

Consider a population of organisms with a set of alleles over the population (the gene pool). Selection will act on members of the population in proportion to the total fitness of each organism to produce a certain statistical distribution of alleles in each of the varying environments the population inhabits. Other effects might distort this distribution, so the selection component will pick out a distribution that has a certain propensity. Given the connection between propensities and probabilities, we can define the fitness of an allele in an environment as the probability of the allele given the environment, where the probability is understood as the propensity due to selection:

$$f_{i,j} = p(a_i|E_j) = \frac{p(a_i \& E_j)}{p(E_i)} \qquad \textbf{(1)}$$

The fitness $f_{i,j}$ of the ith allele in the jth environment is the propensity due to selective forces for the allele a_i to occur in the environment E_j. By probability theory, this is equal to the probability of the allele and the environment both occurring, divided by the probability of the environment. Suppose, for example, we are considering malarial resistance in an environment with malaria present. We know there are three related traits: malarial resistance, no resistance, and the sickle cell trait. We know that the probability of malarial resistance is relatively high, but we also know that, because it is heterozygotic in origin, that there will be a certain amount of the deleterious sickle cell allele. The propensity for this deleterious allele derives from the genetic conditions that produce it. The sickle cell allele is actually somewhat fit, given both the environment and genetic conditions. If we consider an environment in which malaria is not present, then malarial resistance is less fit, and the sickle cell allele is correspondingly less fit. This method of defining fitness allows us to take into consideration interactions between an allele and other factors.

It would be nice to be able to define the fitness of a trait in a similar manner, but we immediately run into a problem. A propensity distribution of traits would contain information not only about the fitness of the traits, but also about the correlation of among traits due to pleiotropic effects. Equation (1) implicitly presumes that the alleles are independent, i.e., that a given allele does not contain information about the probabilities of other alleles. This is not true for traits, since they depend on two alleles (at least). To determine the fitness of traits, we must discount the fitness for the effects of correlation with other fit traits. In the sickle cell case, the fitness of the sickle cell trait can be zero, but it can still appear due to its correlation with the malaria resistant trait. In the following, for convenience, I am going to assume that the discounting for correlations has been done. A similar procedure could be done to correct (1) for gene linkages. For traits, (1) becomes

$$f_{i,j} = p(T_i|E_j) = \frac{p(T_i \& E_j)}{p(E_i)} \qquad \textbf{(1')}$$

Where the selection propensity for trait i in environment j, $p(T_i \& E_j)$, is discounted for pleiotropic effects. The propensity of a trait to appear under selection can be recovered by adding the pleiotropic component back in. The discounting move is required to avoid double counting selection effects. This sort of manoeuvre must be done very carefully to generalise (1') to traits in a full developmental system.

Interestingly, we can sum or even integrate over environments to judge the fitness of a trait under varying conditions (say, for example, the prevalence of malaria varies from location to location). This can be done either for an individual or an interbreeding population to get the fitness of the trait for the individual according to the average time it spends in each environment, or for the population, according to how it is distributed among different environments. The net trait fitness in this case is the weighted sum of the fitness in each environment:

$$f_i = \sum_i p(E_j)\, p(T_i|E_j) = \sum_i p(T_i \& E_j) = p(T_i) \qquad \textbf{(2)}$$

This is just the average discounted propensity of the trait. We could also apply the same method to the fitnesses of individual alleles, depending on their physical and genetic environments.

The adaptive information content of a trait, intuitively, is a measure of how well it matches the environment. On the statistical account of information, it is a function of the logarithm of a probability. Since the fitness is a probability, the natural thing to do is to take its logarithm and take the sum over environments, weighted by the fitness:

$$I_{T_i} = \sum_i f_{i,j} \log_2 f_{i,j} \qquad \textbf{(3)}$$

This, however, would not give us the information of adaptation of a trait, but would give us the information about the distribution of the trait among environments. We need to subtract the information we have about the trait in general, and weight the result by the probability of the environment (in addition to weighting by the fitness) in order to get the average information the trait gives about environments:

$$I_{T_i,E} = \sum_j \mathrm{p}(E_j) f_{ij} \log_2 \frac{f_{ij}}{\mathrm{p}(T_i)} \qquad \textbf{(4)}$$

It is worth noting that the environments do not have to be actual, as long as we has some way to estimate the relevant propensities. I_T, the information in the trait, is maximal for a trait and range of environments in a population when the fitness of the trait in each environment is optimised. If some T is improbable, but highly likely in a common environment nonetheless, it has a high information content, as might be expected. It should also be noted that a trait which is fit in many environments has a higher information content than if it were fit in only one environment, again as we might expect. It is somewhat artificial to use traits in these equations, since traits are not fit in isolation, but traits can be replaced with phenotypes. With the choice of suitable environments, the same equations apply to genic selection (Williams, 1966; Dawkins, 1979) or to group selection (Lewontin, 1970). Equations (1)-(4) follow logically from the propensity definition of fitness, and are no more and no less adequate than that definition. In particular, the definition tells us nothing of the source of the propensities.

We can expand the procedure further to consider a range of traits in a variety of environments. This requires a summation over traits as well:

$$I_{T,E} = \sum_i \sum_j \mathrm{p}(E_j) f_{i,j} \log_2 \frac{f_{i,j}}{\mathrm{p}(T_i)} \qquad \textbf{(5)}$$

Substituting the value of $f_{i,,j}$:

$$I_{T,E} = \sum_i \sum_j \mathrm{p}(T_i \& E_j) \log_2 \frac{\mathrm{p}(T_i \& E_j)}{\mathrm{p}(T_i)\,\mathrm{p}(E_j)} \qquad \textbf{(6)}$$

This is a measure of correlation: The information of adaptation is just the mutual information of alternative environments E and alternative traits T.

The use of propensities, especially the discounted propensities, required to derive these results might be questionable. There is another version of information theory based in computation theory that provides a better justification for the use of discounted propensities. On this account, the information content of a string is the length of the shortest computer program

that can produce that string. By using 1-1 maps of more general forms to strings, we can account for the information content of anything with a definite form. One problem with this approach is that there is a computational overhead involved in the definition. This overhead, though, is small for high information strings, and it largely drops out when we are comparing string lengths. The information of x given y is the information in x, but not in y, i.e. the length of the shortest program to compute x given an input y. The information of adaptation in (6) is a measure of the information that does not need to be computed, given the information in the organism, to determine the environmental conditions. Beyond the limit of this information, the environment will be effectively unpredictable by the organism, and it will be effectively random. This justifies the use of propensities. Justifying the discounting manoeuvre follows a similar pattern: correlations between traits, or linkages among genes, amount to mutual information. As such, even if it is independently embodied in different places, the total amount of environmental prediction that can be based on it is limited to the non-redundant part of the sum of the information of all traits (or genes), so any mutual information should be counted only once.

The information of adaptation is a relation between the environment and the biological entity (gene, trait, organism or group). Since by definition it depends on selection, there is a (somewhat metaphorical) sense in which the biological entity represents the information in its environment. This sense of representation is strengthened by the observation that the origins of differences in phenotypic form within the genealogical hierarchy (Eldredge and Salthe, 1984; Eldredge, 1985) are largely (but not exclusively) genetic. For genetically determined adaptations it is not unreasonable to say that the information of adaptation is encoded in the DNA. Again somewhat metaphorically, the DNA is capable of representing information about the environment. In the same metaphorical sense, DNA can represent information about organisms and species. It is worth noting, however, that equation (6) is completely symmetrical (being a correlation), and does not itself sanction the representational metaphor. We might just as well say that the environment represents the organisms in it. The important thing is that increasing adaptation leads to increasing correlation, which is equivalent to greater mutual information between the environment and the biological entities.

5. Directional Evolution

Natural selection is a powerful tool for explaining the existence of adaptive traits and their functionality. In its neo-Darwinian guise, however, it is deficient for explaining the growth of biological organisation and complexity. A respected biology text (Luria, Gould and Singer, 1981: 647) states "Darwinism is not a theory of intrinsic progress," and goes on to note that adaptations are local only. If that were the whole story, selection could account at best only for local increases of information. Appropriate conditions could just as easily lead to adaptive decreases in organisation, and complexity. Particularly if environmental conditions go through long cycles, we might observe cyclic gain and loss of information. There is no obvious reason increased organisation and complexity should be adaptive, especially throughout large cycles in environmental conditions (Maynard Smith, 1988). Why is selection for apparently simpler forms (as, for example, in viruses, parasites and yeasts[2]) relatively rare? The information of adaptation will increase the most if there is more refined adaptation. A precondition for this is greater variety both among and within organisms. If the intuition that adaptation involves a growth of

[2] Even in these cases there is evidence that this decrease is compensated for by complexity increase across a clad (Brooks and Wiley, 1988).

information is correct, then not only must it be temporally directed, but it must also involve the production of variety.

There are neo-Darwinian accounts of the growth of organisation (Mayr, 1982; Dawkins, 1987: 169-220; perhaps Schanck and Wimsatt, 1988), but these accounts do not arise from the core of neo-Darwinism, if the core is restricted to natural selection as it is usually understood. One approach is to maintain that the organisation produced by adaptation, and the irreversibility of evolution are independent of each other. Irreversibility, on this account, comes from the constant change in the genotype: If a need arises again for some capacity that has been lost, the capacity will be generated by a different genotype than the original (Mayr, 1982: 609). The change in the genotype is partly due to selection, and partly due to random changes (i.e., changes that are not ordered with respect to selection). Although this process can produced increased complexity if the starting point is towards the simpler end (McShea, 1994), there is no clear reason why this process should yield any sort of directionality of the sort that Maynard Smith requires: given two (relatively complex) states of a biological system, there is nothing in Mayr's account of irreversibility that would allow a biologist to tell which state is earlier. Furthermore, since the irreversibility is due to chances events, and not selection, progress is out of the question.

One block to decreasing order is what Wimsatt (Schanck and Wimsatt, 1988) calls the *generative entrenchment*. Once a particular function or structure has been established, functions or structures that evolve later are likely to depend on the earlier evolved traits. Consequently, changes in the earlier traits are less likely to produce a viable organism, since changes to the earlier traits will likely disrupt the traits that depend on them as well. Dawkins (1987: 169ff) pointed out that once a particular chemical pathway has become established, it has an advantage over possible competing pathways, because pathways generally involve interdependencies that must co-evolve together. It is unlikely that a competing pathway will appear, with all the required interdependencies, when an effective pathway is already present. This idea can be generalised to other evolved traits in any circumstance in which co-evolution is required. This mechanism can account for increases in the magnitude of diversity: Once a particular set have traits has co-evolved, any further evolution will tend to continue to diverge in a direction that preserves these traits together. Dawkin's ecological examples involve evolutionary "arms races" that are examples of the "Red Queen" phenomenon, producing directionality, but no net gain in fitness. Together, Wimsatt's generative entrenchment and Dawkins' co-evolutionary fixation of diversity can explain why order does not decrease, and why existing diversity tends to be amplified. They do not, however, explain progress, since they do not require an increase in adaptation.

One way to explain the origins of variety is to postulate isolating mechanisms that divide otherwise unified populations. No doubt these mechanisms are important. External forces (such as geological changes) can separate a population into two or more groups that, no longer reproductively cohesive, will evolve in different directions due either to differing environmental conditions or to chance differences between the genomes of the separated populations (the "founder effect"). Such external influences will need to be invoked by any evolutionary theory, since their presence is well-established. Sexual assortment can serve as an internal mechanism to achieve isolation. It is an interesting question whether there are other internal mechanisms. But again, this mechanism implies no increase in adaptation, and hence no progress.

6. Progressive evolution

The above give us four separate phenomena that can produce directionality: the local phenomenon of adaptation on random variants, mechanisms like the generative entrenchment and

positive feedback that produce irreversibility, mechanisms for the production of variation, and mechanisms to amplify variants. To some extent these overlap and interact, but they are largely independent. This is hardly parsimonious, and reflects the fact that the explanations have been largely added on to the core of traditional neo-Darwinism. It would be nice to have a more systematic explanation. Furthermore, these mechanisms give directionality, but not according to any standard, so they cannot account for evolutionary progress.

One possibility involves entropic increase of information, along the lines described by Brooks and Wiley (1988; Wiley 1988a, 1988b) in the absence of adaptation. If we assume that the environment contains genetic variation, much as a bell jar contains a gas, then we have a source of cohesion that allows us to distinguish genetic macrostates (the environmentally constrained state) from genetic microstates (all of the variant states permitted by a particular environment). This distinction is similar to that I proposed (Collier, 1986) for a macrostate - microstate distinction between species information and the genetic information of all the possible varieties permitted by a given species. If genetic variation approaches equilibrium with the environment, then adaptation will cease. In fact, a strong adaptive force (large fitness differences) will lead to this happening quite quickly. We can expect that near to equilibrium, adaptation will be static. Furthermore, near to equilibrium processes can go in any direction, just like adaptation in traditional neo-Darwinism. But this is not what we observe. We observe a distinct direction of increase in the information of adaptation. This suggests that the genome is *not* in equilibrium with the environment.

One reason to expect non-equilibrium is an analogue to the Brooks-Wiley mechanism for increasing genetic phase space (Collier, 1986). As new variants are selected, the genome becomes more highly adapted, constraining the genetic phase space more severely. On the other hand, the same process often opens up new possibilities for adaptation, which constitutes a release of constraints, possibly making more genetic information potentially adaptive. This would increase the effective genetic phase space. Furthermore, the increased overhead from new possibilities allows the possibility of intermediate levels of organisation that are permitted by non-equilibrium conditions, but not under equilibrium (for example, chance concentrations of particular genes in a relatively small group of closely interbreeding organisms). These spontaneously organised structures might give further chances for adaptation. On a more mundane level, mutations and recombinations continually add information to the genetic phase space which are potential mutual information with the environment. Other mechanisms that can create new possibilities for adaptation, increasing the size of the adaptive phase space are described by Wimsatt and Schanck (1988), Conrad (this volume) , and Baatz (this volume). Unless new forms are maintained strongly and spread rapidly through the population (as would happen if they were especially fit), they are likely to be formed faster than they can equilibrate, thus maintaining non-equilibrium conditions. It is noteworthy that the main reason why equilibrium might be reached reflects exactly the thinking found in traditional neo-Darwinism: perturbation followed by return to equilibrium. However, among all the possibilities, it is an extreme case. These extreme cases would, of course, provide strong confirming evidence for the traditional view, at least in the short run.

Both the production of variability required for unidirectional increases in the information of adaptation, and its unidirectionality itself, can be explained by a genetic system that is well out of adaptive equilibrium with the environment. It is interesting that such conditions enhance the possibilities for adaptation. This might seem counterintuitive, but it is merely the counterpart to the stasis that results from steep peaks in the adaptive landscape. Away from equilibrium, processes involved in adaptation are both more turbulent and more gently productive. These

processes permit not just directional change, but increases in the information of adaptation. Evolution in these regions is the winner in the long run.

References

Ayala, Francisco J. 1988. Can "progress" be defined as a biological concept?, in *Evolutionary Progress* Nitecki (ed), pp 75-96.

Brandon, Robert N. 1990. *Adaptation and Environment*. Princeton: Princeton University Press.

Brooks, Daniel R. and Edward O. Wiley. 1988. *Evolution as Entropy: Toward a Unified Theory of Biology, 2nd edition*. Chicago: University of Chicago Press.

Collier, John D. 1986. Entropy in evolution. *Biology and Philosophy* **1**: 5-24.

Eldredge, Niles. 1985. *The Unfinished Synthesis: Biological Hierarchies and Modern Evolutionary Thought*. New York: Oxford.

Eldredge, Niles and Stanley N. Salthe. 1984. Hierarchy and evolution. *Oxford Studies in Evolutionary Biology* **1**: 184-208.

Griffiths, Paul and Russell Grey. 1994. Developmental systems and evolutionary explanation. *Journal of Philosophy* **91**: 277-304.

Hull, David. 1988. Progress in ideas of progress, in *Evolutionary Progress* Nitecki (ed), pp 27-48.

Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge: Cambridge University Press.

Lewontin, Richard C. 1970. The units of selection. *Annual Review of Ecology and Systematics* **1**: 1-14.

Luria, Salvador. E., Stephen J. Gould, Steven Singer. 1981. *A View of Life*. Menlo Park: Benjamin Cummings.

Maynard Smith, John. 1988. Evolutionary progress and levels of selection, in *Evolutionary Progress*, Nitecki (ed), pp 219-230.

Mayr, Ernst. 1982. *The Growth of Biological Thought: Diversity, Evolution and Inheritance*. Cambridge: Harvard University Press.

McShea, D.W., 1994. Mechanisms of large-scale evolutionary trends. *Evolution* **48**: 1747-1763.

Mills, Susan and John Beatty. 1979. The propensity interpretation of fitness. *Phil. Sci*. **46**: 263-286.

Nitecki, Matthew H. (ed) 1988. *Evolutionary Progress*. Chicago: University of Chicago Press.

Ruse, Michael 1988. Molecules to men: Evolutionary biology and thoughts of progress, in *Evolutionary Progress* Nitecki (ed), pp 97-126.

Salthe, Stanley N. 1993. *Development and Evolution: Complexity and Change in Biology*. Cambridge MA: MIT Press.

Schanck, J.C. and W.C. Wimsatt. 1988. Generative entrenchment and evolution. *PSA 86, Volume 2*, A. Fine and P.K. Machamer (eds.) East Lansing: Philosophy of Science Association, 1988: 33-60

Sober, Elliott. 1984. *The Nature of Selection*. Cambridge, MA: MIT Press.

Weber, Bruce, David Depew, and James D. Smith (eds.) 1988. *Information, Entropy and Evolution*. Cambridge, MA: MIT Press.

Wesson, Robert. 1991. *Beyond Natural Selection*. Cambridge MA: MIT Press.

Wicken, Jeffrey S. 1987. *Evolution, Thermodynamics and Information: Extending the Darwinian Paradigm*. New York: Oxford University Press.

Wiley, Edward O. 1988a. Entropy and evolution, in *Information, Entropy and Evolution* Weber, Depew and Smith (eds), pp 173-188.

Wiley, Edward O. 1988b. Entropy, evolution and progress, in *Evolutionary Progress*, Nitecki (ed), pp 275-292.

Williams, George C. 1966. *Adaptation and Natural Selection*. Princeton University Press.

Wilson, E.O. 1978. *On Human Nature*. Cambridge, MA: Harvard University Press.

Wimsatt, William C. and Jeffery C. Schanck. 1988. Adaptations and the means of their avoidance, in *Evolutionary Progress*, Nitecki (ed), pp 231-274.

Wright, Larry. 1973. Functions. *Philosophical Review* **82**: 139-168.