

Mutualism, Parasitism, and Evolutionary Adaptation

Richard A. Watson¹

Torsten Reil²

Jordan B. Pollack¹

¹ Dynamical and Evolutionary Machine Organization

Volen Center for Complex Systems – Brandeis University – Waltham, MA – USA

² Department of Zoology – University of Oxford – Oxford – UK.

richardw@cs.brandeis.edu

Abstract

Our investigations concern the role of symbiosis as an enabling mechanism in evolutionary adaptation. Previous work has illustrated how the formation of mutualist groups can guide genetic variation so as to enable the evolution of ultimately independent organisms that would otherwise be unobtainable. The new experiments reported here show that this effect applies not just in genetically related organisms but may also occur from symbiosis between distinct species. In addition, a new detail is revealed: when the symbiotic group members are drawn from two separate species only one of these species achieves eventual independence and the other remains parasitic. It is nonetheless the case that this second species, formerly mutualistic, was critical in enabling the independence of the first. We offer a biological example that is suggestive of the effect and discuss the implications for evolving complex organisms, natural and artificial.

1 Introduction

The phrase “survival of the fittest”, ubiquitous in our thoughts about evolution, is often taken to mean mutually exclusive competition. Accordingly, mutually beneficial relationships are generally treated as a curio. But biological evidence suggests that mutualism is an important enabling mechanism in evolutionary innovation. In its strongest form, symbiosis can lead to *symbiogenesis*: the genesis of new species via the genetic integration of symbionts [Khakhina 1992, Kozo-Polyansky 1921, Margulis 1992, Merezhkovsky 1909]. For example, eukaryotes, which include all plants and animals, have a symbiogenic origin [Margulis 1992].

Such ‘genetic integration’ may occur via direct genetic mechanisms such as horizontal gene-transfer, but our earlier work [Watson & Pollack 1999a] provided a simple model of a relatively subtle, indirect mechanism whereby the genetic characteristics of one organism may be acquired by another. Our model parallels Hinton and Nowlan’s work, “How Learning Can Guide Evolution” [1987]. Their paper demonstrates the Baldwin effect [Baldwin 1896], a phenomenon whereby learned, or plastic, characteristics can induce equivalent innate, or non-plastic, characteristics. In our adaptation of their model we replaced learning with symbiosis; or more generally, replaced the lifetime plasticity of an organism

with lifetime interaction between organisms. This enabled us to show how the characteristics of one organism can be induced in another symbiotic organism. This shaping effect enables the evolution of organisms that would otherwise be unobtainable—or at least, would be very unlikely to occur.

Non-genetic variation guides genetic variation

Our simulation of the effect can be described in two phases. First, symbiotic groups find the solution to a problem (a set of abilities that confers high reproductive fitness) more quickly than the solution can be found by a single organism. This occurs simply because the combinations of abilities formed via lifetime interaction of organisms samples a much larger set of variations than the relatively slow genetic variation from mutation. In the second phase, after an ecosystem of mutually beneficial organisms has become established, the evolution of the individual organisms therein operates in a different environment. Where previously an organism that exhibited some fraction of the necessary abilities, but not all the necessary abilities, would fail, now symbionts may *fill-in* for this organism’s inadequacies. Moreover, the greater the fraction of necessary abilities it exhibits the less filling-in is required—i.e. the less it depends on symbionts and the more reliably successful it is. This provides a gradient to guide genetic search toward an organism that can ultimately perform independently. Without the support of symbionts this gradient does not arise and therefore the occurrence of an independent organism exhibiting the solution would require an improbably fortunate random mutation.

Thus, the abilities first discovered by the symbiotic group become encapsulated in the heritable traits of a single individual. We call this effect *symbiotic scaffolding*: the symbionts support each other as partially able organisms, and enable the gradual accumulation of abilities, until ultimately, when their abilities are complete, the scaffolding is not required.

One way to interpret how this effect operates is as a ‘smoothing’ of the fitness landscape. The dotted curve in Figure 1 shows an arbitrary rugged fitness landscape. Each point on the horizontal axis represents a phenotype (or set of traits), and the dotted curve indicates the fitness of the phenotypes when evaluated independently. Now imagine that when an organism is evaluated, its own characteristics may be supplemented with those of other

organisms. This modified set of characteristics will still contain the characteristics of the original organism and will therefore be somewhere in the neighborhood of the original organism's phenotype. If the organism interacts with many other organisms during its lifetime then its fitness will reflect the fitness of a set of points sampled from the neighborhood of its own characteristics. Points on the solid curve in Fig. 1 are an average of the original landscape weighted using a Gaussian. This averaging over the local neighborhood acts so as to 'smooth' the fitness landscape and provide a kind of 'look-ahead' about phenotypes in the nearby vicinity. This modified fitness landscape enables genetic search to escape from local optima in the original landscape and move towards fitness peaks that were formerly unobtainable.¹

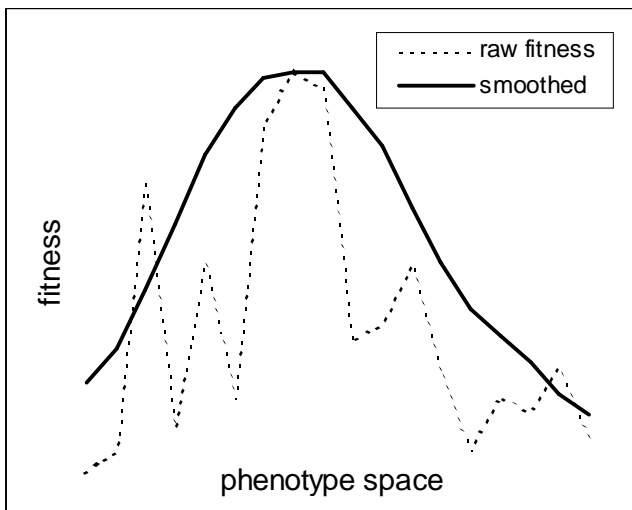


Fig. 1. The guiding effect of symbiotic interaction can be interpreted as a smoothing of the fitness landscape.

Previous experiments with one population

In our previous work we illustrated this effect using an adaptation of Hinton and Nowlan's model of the Baldwin effect. We evaluated individuals in the context of many randomly constructed groups of individuals and, instead of assigning fitness based on their individual ability, we gave them a fitness that reflected the average success of groups they formed. Without the group evaluation individuals could not be evolved to solve the problem. But with the benefit of group evaluation the symbiotic scaffolding effect enabled the evolution of initially mutualist groups that solved the problem together, and ultimately independent organisms that solved the whole problem by themselves. With this effect, we illustrated a

¹ This interpretation of scaffolding is quite similar to Hinton and Nowlan's explanation of the Baldwin effect, as the reader may recognize—especially, when we apply the smoothing to Hinton and Nowlan's problem in Figure 2 (which we will introduce shortly).

mechanism whereby the formation of mutualist groups enabled the evolution of organisms that would not otherwise have occurred. In so doing, this effect also shows how the characteristics of an organism can be induced in another symbiotic organism without direct transfer of genetic material. These are important principles for understanding the role of symbiosis in evolutionary adaptation.

Our model does not use the exchange of genes between organisms—the symbionts may be distinct species—and our original experiments did show some evidence for genetic divergence in the population. However, we used only one population of organisms and thus, in general, the symbionts may have been closely genetically related (by inheritance). This weakened the interpretation of the model as symbiosis between different species, and brought into question whether the effect would be seen in an ecosystem of truly separate species.

Mutualism between separate species

In the new experiments we present in this paper we sought to verify that the phenomenon was reproducible when the symbionts were reproductively isolated species. This would show the effect to be relevant to symbiosis between genetically unrelated organisms and therefore more widely applicable in nature.

Thus, in this paper, we form groups by selecting individuals from two separate fixed-size populations each of which reproduces independently. (In all other respects the experimental setup we describe here is the same as that used in our earlier experiments.) As expected, the scaffolding effect is still observed; confirming our hypothesis that mutualist groups consisting of genetically separate species can be instrumental in catalyzing the evolution of complex independent organisms.

But there is an interesting difference from our earlier results. When the symbiotic group members are drawn from two separate species, only one of these species achieves eventual independence. The pressure for the second to become independent falls off and it becomes a parasite—it gains benefit from its perfect partner but provides nothing in return.

The remaining sections of this paper are organized as follows. In the next section we describe our experimental setup, and Section 3 gives results. Section 4 introduces an example from nature that is suggestive of symbiotic scaffolding. Finally, we discuss the implications for artificial life research, and conclude.

2 Experimental set-up

Hinton and Nowlan provide a simple and elegant abstract model of the Baldwin effect which has been replicated and extended many times [Belew 1989, Harvey 1993, Mayley 1996]. Our experiments use an adaptation of their 'extreme and simple scenario'. The model is deliberately abstract so that the combinatorics involved in the effect

are clear. We consider a problem that consists of a large number of variables all of which must be correctly specified by an organism in order for that organism to receive any reproductive fitness. In such cases an organism that is partially correct, even one that specifies all but one of the variables correctly, is not rewarded at all. This worst-case scenario is the extreme case of irreducible complexity, in which solutions can only be found by trying possibilities at random.

As an example scenario, we may imagine a metabolic chemical cycle with 20 steps. Each of the 20 steps must be performed by an organism correctly in order to get the chemical cycle going and to thereby confer reproductive fitness.

As stated thus far, this ‘needle on a plateau’ fitness landscape provides no gradient to lead search towards the solution – an organism with 19 correct steps is not favored over an organism with, say, only 1 correct step. But when we introduce lifetime interaction between organisms this will enable a gradient that leads genetic variation toward the solution. The smoothing afforded by group interaction modifies the fitness of points near the solution so that they are preferred over points farther from it, as depicted in Fig. 2.




Fig. 2. An illustration of the simple problem space used in the following experiments (and in those of Hinton and Nowlan), depicting the ‘smoothed’ landscape effected by symbiotic interaction.

Evaluating Groups

To model lifetime interaction we test groups of organisms instead of individual organisms, and we test an organism in many groups during one lifetime. Each organism may prohibit, enable, or have a neutral effect on a step in the chemical process. If an organism is neutral with respect to a step then this step may (or may not) be completed by some other organism (of a different species). That is, an organism can gain the benefit (or penalty) of chemical byproducts created by the processes of other organisms in

the ecosystem (for those steps where the organism itself is neutral).

We may crudely represent the relevant traits of an organism in 20 genes where each gene has three alleles: *correct*, *incorrect* or *neutral* corresponding to enabling, prohibitive (preventing completion of the cycle), or neutral interaction with a step in the cycle. This unrealistic simplification enables us to see the mechanisms of interest more clearly but it is not integral to the results that follow.

We use the happenstance co-location of organisms to form groups since it makes minimal assumptions about the nature of symbiotic interactions. The use of a more sophisticated model of symbiotic relationship-forming will illustrate the scaffolding effect more strongly—we stress that our model of organism interaction is deliberately trivial so as to prevent details from obscuring the essence of the effect. In our model we may imagine that organisms are randomly distributed in the environment and perpetually mixed. At any one instant there will be some number of other organisms in the immediate vicinity of the organism in question. Thus every organism is tested by combining its abilities with those of several other randomly selected organisms of the other species. Fig. 3 shows how the abilities of organisms are combined. In this figure there are two species, A and B, and we prevent solutions being provided by mutualism within a species by using only organisms from the complementary species to fill-in groups.

first organism, A1:	00-0-11---01-1--0--1
second organism, B1:	-1-0-1-001-01-1--0--
third organism, B2:	1--10-11-0----10---1
fourth organism, B3:	-01---0-0----10-0-11
fifth organism, B4:	--111---0101--0-0--
combined abilities:	00100110010111100011

Fig. 3. Combining the abilities of organisms. The 20 traits of each organism may take one of three forms: correct, incorrect or neutral shown as 1, 0 and “-”, respectively. The neutral abilities of an organism from species A may be filled-in by the abilities of organisms from species B. Notice that the traits of the first organism take priority over all others; for consistency, the traits of the second organism take priority over all but the first, and so on. In some groups, some organisms may be redundant since every trait is specified by at least one of the preceding organisms, as illustrated by the fifth organism shown here. Note that this combination of abilities does not represent the formation of a new organism - it is merely a representation for the result of different species acting in concert.

Since both the selection and ordering of the organisms are random, the details of this filling-in mechanism are largely inconsequential to the results that follow. One important feature, however, is that the fitness of the combined traits is awarded to the first organism only, and that the traits of the first organism are not over-ruled by

any other. However, since the first organism will likely fill-in for other organisms in their turn, this asymmetry is reciprocated. Alternate models of interaction and reward distribution may be equally valid; however, the current model is sufficient for our purposes.

A key feature of the mechanism we are modeling is that the search for combinations of abilities via lifetime interaction is much more rapid than that arising from genetic variation alone. Hence we test each organism in 1000 random groups during its lifetime. Each group is formed from a different random selection of organisms drawn from the other species.

Our earlier experiments showed that when appropriate symbionts are reliably available, and incur no additional overheads, then there is no pressure to be independent. But naturally, if there is some cost to relying on symbionts then independent organisms are preferred. In our earlier experiments we added an implicit cost by limiting the availability of symbionts. Implementationally, we limit group sizes probabilistically with the limit randomly selected from an exponential distribution for each group formed. Specifically, the probability of there being exactly k members in a group is 2^{-k} , $k \geq 1$. In this way it is most likely that an organism will be evaluated on its own; next most likely it will be evaluated with one other organism, and so on.

Finally, the fitness of an organism is given by $f=1+n$, where n is the number of groups (out of the 1000 groups tested) in which the organism in question forms a successful group.

Experiments

The genetic model, the method of interaction, and the evaluation described above are iterated in a genetic algorithm (GA) [Holland 1975]. Hinton and Nowlan chose the population size, number of lifetime trials, and number of variables in the problem carefully so as to make it most unlikely that genetic variation alone would find the solution but very likely that lifetime variation would. We continue to follow the experimental parameters of Hinton and Nowlan where applicable for the same reasons. We use 1000 individuals but here they are divided into two populations of 500 representing two species. Each species, say **A** and **B**, reproduces independently, and as indicated in Figure 3 (above), each member of **A** is evaluated by filling-in its missing abilities with members from **B**, and vice versa. When an individual is being evaluated the group will not contain other members of that individuals' species.

Fitness-proportionate reproduction is applied generationally [Holland 1975] to each population separately. In this way there is no competition between the members of population **A** and the members of population **B**, but the specific configuration of individuals in each population does have an effect on the fitness of individuals in the other population through the group evaluation in which they are involved.

In these experiments we use mutation as our only

source of genetic variation. Mutation is applied with a bitwise probability of 0.05 of assigning a new random value. New values are randomly selected to be correct, incorrect or neutral genes with probability 0.25, 0.25, 0.5 respectively. These same proportions are used to construct the initial population so that, on average, an individual will have half its genes neutral and half non-neutral (as illustrated in Fig. 3).

We should emphasize which parts of the results that follow are expected and which parts offer new insight. As stated, the size of the problem (number of steps in the cycle), ratio of alleles in the initial population, number of organisms, and number of groups tested for each organism are deliberately chosen to make the discovery of a successful individual most unlikely and the discovery of a successful group most likely. The interesting part of the effect is what happens *after* successful groups are formed. Specifically, we are looking for how the trends in the make-up of individuals change in the context of symbiotic groups. We will see that before the formation of successful groups, increases in the number of correct alleles are not selected for. However, after groups are formed, there is a trend towards more correct alleles. This trend eventually results in some organisms becoming independent which would not have happened without the presence of the (eventually redundant) mutualists.

3 Results

Fig. 4 shows the number of each allele per organism averaged over all organisms (in both populations) at each generation. We see that the proportion of alleles at the start of the experiment is as per the mutation probabilities, i.e. approximately 0.25,0.25,0.5 for correct, incorrect and neutral respectively. Around the 225th generation a quite dramatic change takes place: the proportion of incorrect alleles falls close to zero whilst the number of correct alleles rises. (The exact generation at which these sharp changes occur varies from run to run due to the stochastic nature of the experiment.) This is the point where symbiotic organisms become established and incorrect alleles are purged from the gene pool. Thereafter we see a clear upward trend in the number of correct alleles in subsequent generations. Unlike our original experiments there is no clear trend towards 20 correct alleles in the following generations. But we will see in a moment that, although the average number of correct alleles over both populations does not continue to rise, some individuals do find all 20 correct alleles.

In Fig. 5 we see that the dramatic changes around 225 generations coincide with the establishment of groups that solve the complete cycle.² Then we see, at about 350 generations, the establishment of individuals that are self-

² Although there are a few instances of successful groups in the first 200 generations they do not take hold in the population.

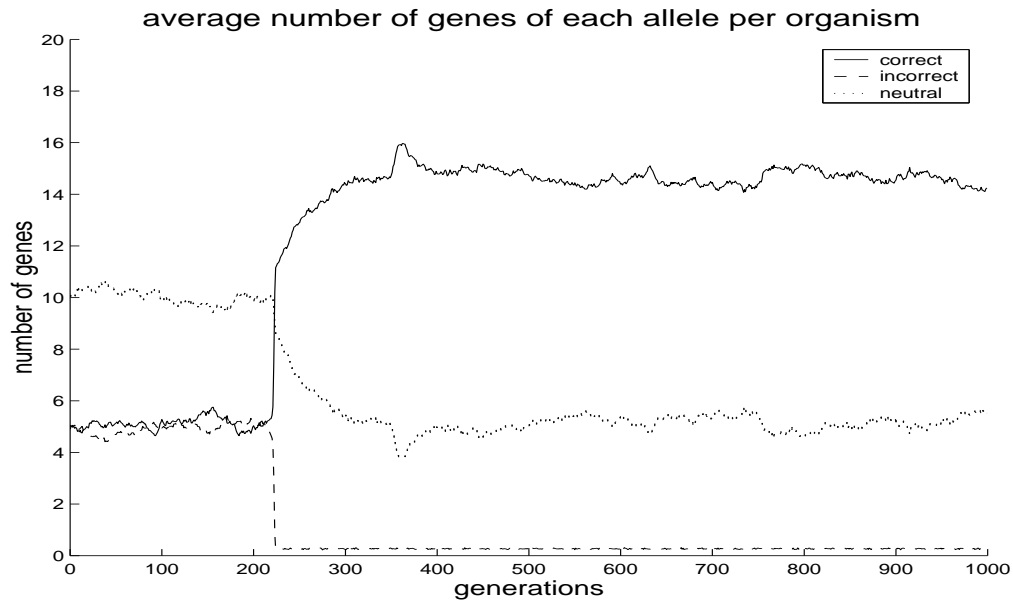


Fig. 4. Number of genes of each allele per organism, averaged over all 1000 organisms at each generation.

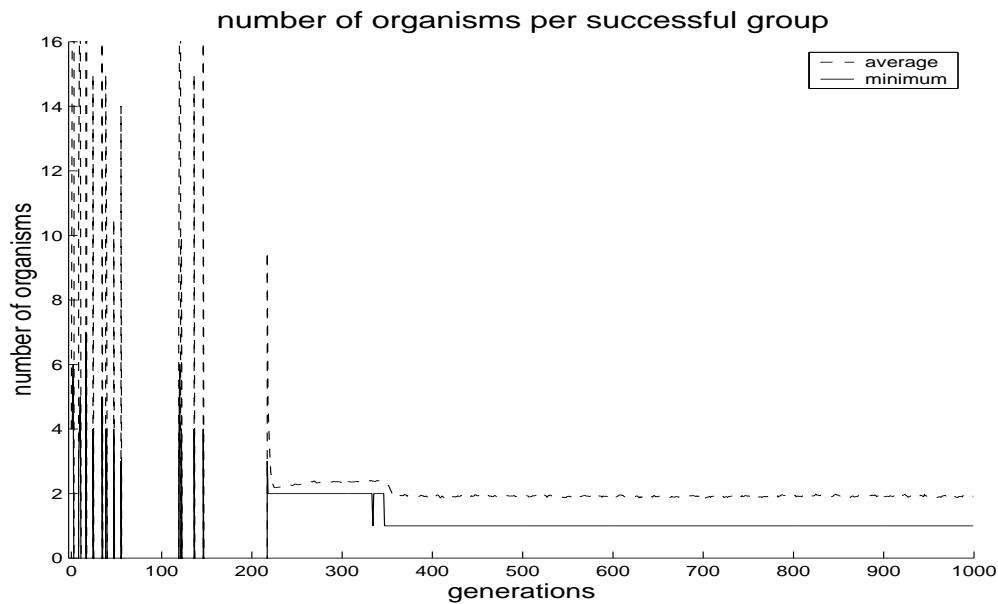


Fig. 5. Average and minimum size of successful groups. Where no successful group is formed (as in some of the first few generations) both the average and minimum group sizes are shown as zero.

sufficient (i.e. we see a minimum group-size of 1).

The effects shown in Figures 4 and 5 have some qualitative similarity with our previous work using a single population. This is as expected. Separating the organisms into two reproductively isolated populations makes no difference to the combinatorics and probabilities involved, and the establishment of complementary mutualists does not depend on genetic relatedness. However, we need to explain why the average number of correct alleles remains lower than in our single

population model, and when we examine the make-up of the two populations separately in Fig. 6 (and also Fig. 7) some interesting features are revealed.

From Fig. 6 we can see that in the initial stages the maximum number of correct alleles shows no significant trends but begins to increase when mutualist groups are established. The trends in both populations are identical thus far. But, as one of the species approaches independence the other species starts to decline. There is no asymmetry between the two populations in the set-up

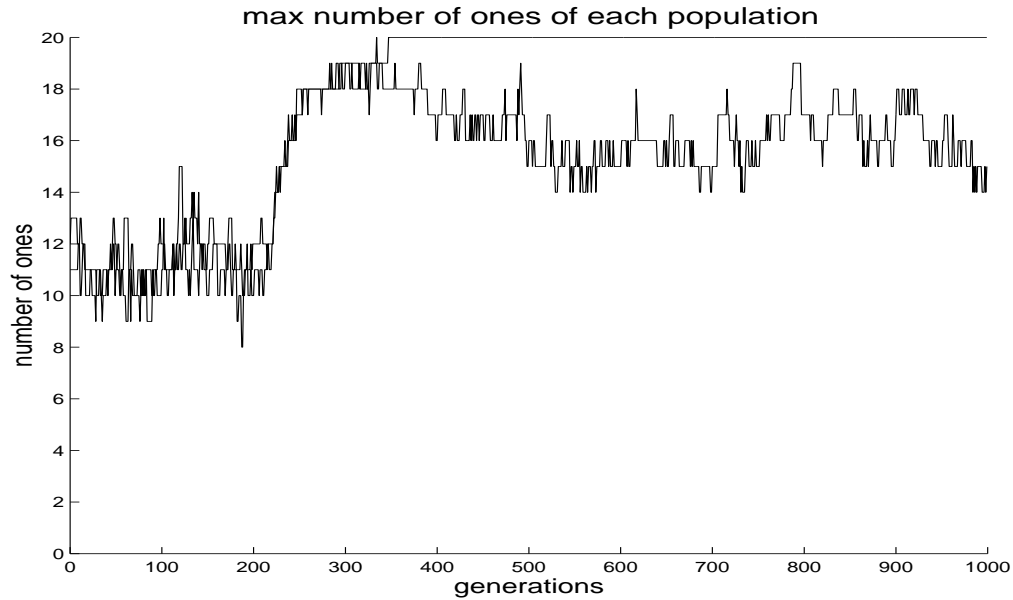


Fig. 6. Maximum number of correct alleles shown separately for each population.

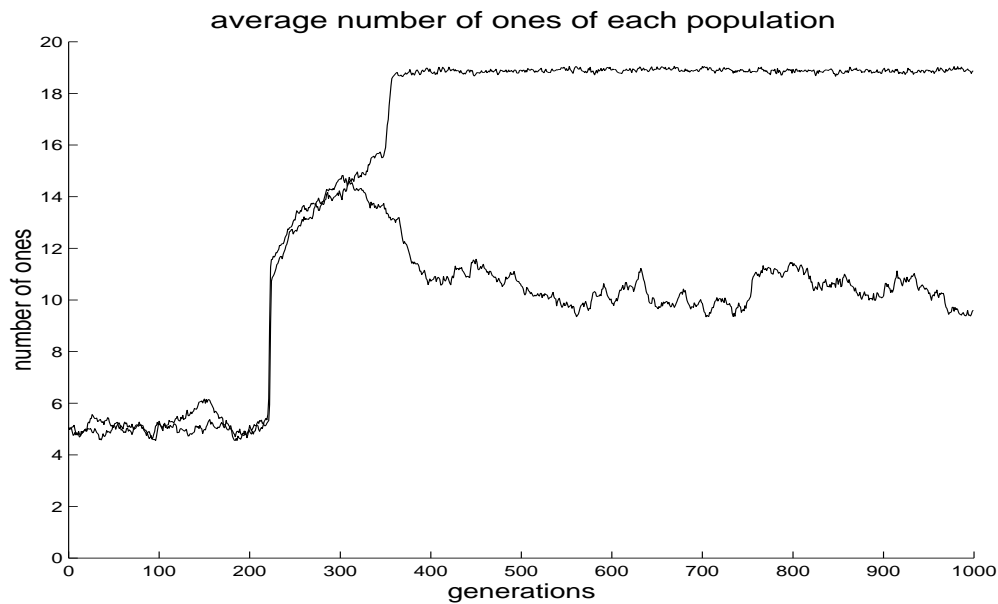


Fig. 7. Average number of correct alleles shown separately for each population.

of the experiment – the asymmetry that develops depends on which of the two populations happens to find independent organisms first. The maximum number of correct alleles in the second population now starts to decline. These trends are seen more clearly in Figure 7 which shows the average number of correct alleles. There is a clear intermediate stage where organisms in both populations are increasing in ability and a clear separation in the latter stage where one species becomes independent. The second species is not contributing anything to the first but does receive benefit from it – it has transformed from a mutualist to a parasite. The rise of

this parasitic species can be understood as follows. Consider two species **A** and **B**. When there are imperfections in members of **A** there is some advantage in increased independence of **B**. But if **A** is perfect then members of **B** with say, one or ten neutral abilities are equally likely to be correctly supported by **A**. Thus, there is no fitness differential for organisms in **B**. The only exception being the preference for an organism that does not need **A** at all. In fact, in a context of perfect hosts, organisms are faced with a ‘needle on a plateau’ problem space that is the same as the original space without symbionts. So, as one species approaches independence

and thereby becomes a nearly perfect host, the pressure for the second species to exhibit correct bits declines.

The exact changes in the second species will be determined by the balance between the remaining selection pressure for correct bits (since the hosts are not quite perfect, on average), and the pressure of mutation which, when individuals have a high number of correct bits, will usually reduce the number of correct bits.

Since the introduction of new alleles via mutation includes both neutral and correct bits³ in the ratio 2:1 (or 0.5:0.25) respectively, we might expect that the population will likely drift to reflect this ratio, i.e. show an average of approximately 33% correct bits. In practice, it seems that the ratio is reliably higher—not going outside 40-60% in all 17 (of 20) runs of the experiment that exhibited a host/parasite split. The reason for this is not clear to us but may possibly reflect the fact that there is still some slight pressure for correct bits in the parasites since the hosts are not quite perfect (about 19 correct bits on average).

Note that the fitness of the parasites is considerably lower than the fitness of the hosts—recall that most of the time an individual is evaluated on its own given the probabilistic group-size of 2^k , and thus, in most trials, a parasite will get a fitness of zero. In contrast, a perfect host gets maximum fitness in all trials. However, since the hosts and parasites are in separate populations the parasites will not be replaced by hosts (as they were in our original experiments). And, unless some member of the parasite population by chance makes the transition to full independence, the fitness of all parasites is the same: when they are evaluated alone they receive no fitness contribution, when they are evaluated with a perfect host they receive full-fitness. Their average fitness in the context of perfect hosts is only a function of the probability that they are evaluated in a group, and has no dependence on the number of correct traits they exhibit.

In these experiments we have focused on models of two species. Preliminary experiments with three or more species indicate that the likelihood of a species becoming independent decreases with the number of other species that are already independent. The last species finds itself drifting in the same manner as the second species in these experiments.

To summarize, in this overall effect we see an ecosystem of unrelated organisms that become mutualist organisms, scaffolding for each others inabilities. During this mutualist period the organisms increase in their individual ability until one of the species becomes independent. The characteristics formerly exhibited only by mutualist groups have been encapsulated in the characteristics of single organisms. As the first species approaches independence the second species transitions from being a mutualist to being a parasite. Although it now confers no benefit to its independent host, it was

essential in enabling the evolution of the host's independence.

In other runs of the experiment there is considerable variation in the exact generation in which symbiotic organisms become established, and in the generation which exhibits the first self-sufficient individual. Nevertheless, the discovery of an independent species and a residual parasitic species is quite reliable. 17 out of 20 runs of the experiment resulted in an independent species with a parasite. As stated, the average number of correct alleles in the second species (through 1000 generations) remained around 50% in these 17 runs. In the other three runs the second species also found independent organisms. In most runs the intermediate stage, where mutualism is increasing the number of correct alleles, occurs very quickly, but without the mutualism (i.e. without group evaluation) the effect does not occur at all, and an independent organism is never found.

4 An example from nature

The search for examples of symbiont scaffolding in the natural world is complicated by the transient nature of the phenomenon. Evidence supporting our original, one-species model was particularly problematic since the end result expected to find nothing but independent organisms. However, in this two-species model we expect a residual symbiotic partner to be more common and this type of relationship is easier to find.

For the process to be identified as such, a symbiotic relationship has to be recognised in which at least one partner has acquired at least one trait that was previously exclusive to the other partner. There must furthermore be evidence that this particular trait would have been unavailable for the partner in question through conventional evolution. We suggest here that these conditions are met by the bivalve species *Solemya reidi*.

Members of the genus *Solemya* are characterised by a close symbiotic relationship with sulfide oxidising bacteria [Cavanaugh 1994] enabling them to inhabit sulfur rich environments such as deep thermal vents, which would otherwise be inhospitable due to the high toxicity of sulfur [Grieshaber & Völkel 1998]. In addition to providing a detoxification mechanism the bacteria supply the clam with organic compounds from CO₂ fixation.

Intriguingly, one member of the genus, *S. reidi*, exhibits the ability to oxidise sulfide by itself – a trait previously unique to its symbiotic bacteria, which it still harbours in its gills [Powell & Somero 1986]. We suggest that this partnership represents an example of symbiont scaffolding as we have modelled. The scaffolding process may be complete, in which case the bacteria are now parasitic. Alternatively, *S. reidi* may represent a transitory stage of the process, that is, the clam has acquired one characteristic from the scaffolding bacteria but the bacteria still provide other characteristics essential to the clam. Which of these scenarios is the case is not yet clear.

The possibility of a transitory condition is supported by

³ Incorrect bits are still selected against.

the fact that *S. reidi* lives in anthropogenic habitats, namely sewage outfalls and pulp mill effluents, making it likely to be an evolutionary novelty. As such, there may have been insufficient time for symbiont scaffolding to come to an end, i.e. to reach a point where the bacteria are no longer of benefit to the clam and are therefore entirely parasitic.

In summary, we propose that the case of *S. reidi* illustrates symbiotic scaffolding as follows: a) The toxicity of the sulfur-rich environments presents an impossibly hard transition for the adaptation of the clam via conventional evolution; the correct adaptation being analogous to our 'needle on a plateau'. b) The mutualist clam-bacteria symbiosis, as exhibited by other *Solemya* species, was an ancestral stage for *S. reidi* before its own ability to oxidise sulfur. c) The provision of sulfide oxidation by the bacteria smoothed the problem landscape presented to the clam. Given some cost in reliance on symbionts, or some benefit in independent sulfur oxidation, this relationship induced *S. reidi* to gradually evolve the trait itself. In this case, the main adaptive advantage in being less reliant on the bacteria was energy production (sulfide oxidation in *S. reidi* is coupled to ATP synthesis in the mitochondria). The extent of the latter is demonstrated by the absence of both mouth and gut in *S. reidi* [Cavanaugh 1994], indicating total independence from conventional feeding mechanisms. d) Thus, the clam has acquired a characteristic from the bacteria via symbiotic scaffolding. *S. reidi* is now less dependent on the bacteria that catalysed its own abilities, at least in respect of sulfide oxidation.

While the case of *S. reidi* is suggestive, one caveat must be pointed out. It is possible that the clam's ability to oxidise sulfide has arisen as a consequence of gene transfer from the bacteria to the bivalve's mitochondria. If such inter-species gene transfer has occurred then symbiont scaffolding, as described in the previous sections, would not be required to explain the characteristics of *S. reidi*. In order to clarify the issue, the sequence of the sulfide oxidising enzyme in both symbiotic partners must be determined and tested for homology.

5 Discussion and Future Work

In symbiogenesis, the creation of new species via the genetic integration of pre-adapted symbionts provides a different source of innovation from the Darwinian gradual accumulation of random mutations. And the indirect acquisition of symbiont characteristics that we have modeled in this paper also suggests a different perspective on adaptation. The scaffolding effect provides a mechanism by which non-genetic variation can guide genetic variation. Although in this indirect mechanism it is still the accumulation of random mutations that implements the effect, the mutations now occur under selection which is educated by a form of 'look ahead'. This look ahead is provided by the exploration of

phenotypic characteristics afforded by the formation of symbiotic groups that include the organism in question.

Though there are many computational models that explore the evolution of cooperative behavior there are none that use mutualism as an integral part of adaptive innovation. As Artificial Life (ALife) researchers we are well advised to understand the sources of innovation in natural evolution. Changes in the way we look at organism interaction, mutualism, and evolutionary adaptation in general, open new directions for ALife research.

For example, existing ALife models that are based on genetic evolution utilize some fixed representation of genes, of one kind or another, and evolutionary search uses fixed variation operators to explore the space this representation affords. A failing that is therefore common to all such systems is their inevitable *complexity ceiling*. That is, after some initial promise, further innovation is not forthcoming since most variations are detrimental [Kauffman 1993]. In light of this, perhaps the most interesting interpretation of the symbiogenic and symbiont scaffolding mechanisms is that they provide natural evolution with an escape from this problem by creating new 'units of variation'. In the models we have illustrated, the important discovery takes place by shuffling combinations of *organisms*, not by mutation of the genes. Genetic mutation merely follows in its footsteps. At first, sets of abilities are explored by shuffling groups of simple organisms. Then successful groups are encapsulated into composite individuals that exhibit the characteristics formerly exhibited by the group. Now sets of abilities may be explored by shuffling groups of these more complex organisms – variation is now operating on larger units. The potential for the mechanisms to recurse in this way provides the opportunity to scale-up the representation in which search takes place. This, we believe, has potential for over-coming complexity ceilings in ALife models, just as it has been instrumental in enabling major transitions in natural evolution [Maynard-Smith & Szathmary 1995]. Other ongoing research is directed at hierarchical problem solving using these ideas [Watson et al 1998, Watson & Pollack 1999b].

6 Conclusions

This paper has developed previous work investigating the guiding effect of mutualism in evolutionary adaptation. This effect, which we call symbiotic scaffolding, enables the characteristics of mutualist groups to become encapsulated in a single individual and thereby enable the evolution of independent complex organisms that would not otherwise occur. We have shown here that this effect may occur between genetically unrelated species and the new set-up has revealed an interesting new feature. When mutualists of two different species scaffold one-another only one achieves eventual independence and the other remains as a parasite. Finally, we introduced a biological example that is suggestive of this model.

Acknowledgments

Though several audiences suggested re-implementing our previous work with separate species, Peter Todd, of the Max Planck Gesellschaft, suggested directly that when the two populations were separated they may exhibit some interesting asymmetry – which indeed they did. Thanks also to Steven Frank and Vladimir Kvasnicka for sharing comments. Thanks, as always, to the members of DEMO at Brandeis for assisting us in the research process.

References

- Baldwin JM, 1896, "A New Factor in Evolution," *American Naturalist*, 30, 441-451.
- Belew RK, 1989, "When Both Individuals and Populations Search," in Schaffer JD, ed., *ICGA3*, San Mateo California: Morgan Kaufmann.
- Cavanaugh, CM, 1994, Microbial Symbiosis: Patterns of Diversity in the Marine Environment, *Amer. Zool.*, 34: 79-89.
- Grieshaber, MK, Völkel, S, 1998, Animal Adaptations for Tolerance and Exploitation of Poisonous Sulfide, *Annu. Rev. Physiol.*, 60: 33-53.
- Hinton GE, & Nowlan SJ, 1987, "How Learning Can Guide Evolution," *Complex Systems*, 1, 495-502.
- Harvey I, 1993, "The Puzzle of the Persistent Question Marks: a Case Study of Genetic Drift," in S. Forrest (ed.), *Genetic Algorithms: Proc. of the Fifth Intl. Conference*, pp. 15-22, San Mateo CA, Morgan Kaufmann.
- Holland JH, 1975 *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: The University of Michigan Press.
- Kauffman, S, 1993, *The Origins of Order*, Oxford University Press.
- Khakhina LN, 1992 *Concepts of Symbiogenesis: Historical and Critical Study of the Research of Russian botanists*, eds. Margulis L, McMenamin M, (translated by Merkel S, Coalston R), Yale University Press.
- Kozo-Polyansky BM, 1921, "The Theory of Symbiogenesis and Pangenesis, Provisional Hypothesis" in *Journal of the first All-Russian Congress of Russian Botanists* in Petrograd in 1921. Petrograd (in Russian, see Khakhina 1992).
- Mayley G, 1996, "Landscapes, Learning Costs and Genetic Assimilation," in *Evolution, Learning, and Instinct: 100 Years of the Baldwin Effect, special issue of Evolutionary Computation*, Vol 4, No. 3; P.Turney, D. Whitley and R. Anderson (eds).
- Maynard-Smith, JM & Szathmary, E, 1995, "*The Major Transitions in Evolution*", WH Freeman and Co.
- Margulis L, 1992, *Symbiosis in Cell Evolution: Microbial Communities in the Archean and Proterozoic Eons*. 2d ed., W.H. Freeman, New York.
- Merezhkovsky KS, 1909 "The Theory of Two Plasms as the Basis of Symbiogenesis, a New Study or the Origins of Organisms," *Proceedings of the Studies of the Imperial Kazan University*, Publishing Office of the Imperial University. (In Russian, see Khakhina 1992).
- Powell, MA, and Somero, GN, 1986, Hydrogen-Sulfide Oxidation is Coupled to Oxidative-Phosphorilation in Mitochondria of *Solemya reidi*, *Science*, 233: 563-566.
- Watson RA, Hornby GS, & Pollack JB, 1998, "Modeling Building-Block Interdependency," *PPSN V*, Eds. Eiben, Back, Schoenauer, Schweffel: Springer.
- Watson, RA & Pollack, JB, 1999a, "How Symbiosis Can Guide Evolution". Procs. of Fifth European Conference on Artificial Life, Floreano, D, Nicoud, JD, Mondada, F, eds., Springer.
- Watson, RA, & Pollack, JB, 1999b, "Incremental Commitment in Genetic Algorithms," In Banzhaf, W, Daida, J, Eiben, AE, Garzon, MH, Honavar, V, Jakiela, M, & Smith, RE eds. *GECCO-99*. San Francisco, CA: Morgan Kaufmann.