A study of the effects of group formation on evolutionary search

MSc Dissertation by
Richard Watson

in partial fulfilment of masters degree in
Knowledge Based Systems

University of Sussex
Falmer, Brighton, U.K.

July 1996

A study of the effects of group formation on evolutionary search

Abstract

The formation of mutually beneficial groups in an ecosystem may enable the creation of complex entities that would otherwise be unevolvable. This process, symbiogenesis,[1] is believed to have enabled certain of the "major transitions in evolution", [Maynard-Smith & Szathmary, 1995].

This work describes several computational models which investigate the mechanisms involved in enabling symbiogenesis. The experiments attempt to establish both how pervasive the influence of mutually beneficial groups is in nature, and the applicability of these mechanisms to practical engineering problems through Artificial Evolution.

The models include a 'Group Effect' which shows how the presence of mutually beneficial groups in an ecosystem can alter the fitness landscape and thereby enable the evolution of more complex species. Another model demonstrates how this Group Effect may enable the re-combination of 'building-blocks' in a manner which provides a real computational advantage over genetic re-combination in the standard genetic algorithm.

Such mechanisms may enable significant increases in the level of complexity that can be attained in evolution, both natural and artificial.

Keywords: symbiogenesis, hypercycles, diversity, cooperation, coevolution, population dynamics, group selection, genetic algorithms (esp. combinatorics of), parasites, hitch-hikers, encoding schemata, building-blocks, Royal Road functions.

Acknowledgments

Many thanks to Inman Harvey for his guidance and patient tolerance, to Dave Cliff for his encouragement in the early stages of establishing a research direction, to the members of the ALERGIC reading group for the many discussions which have assisted in formulating my ideas, and to Jason Noble for his comments on the final draft.

[1] [Margulis 1993] and [Merezhkovsky 1920].

Contents

Figures

Tables

1. Introduction

1.1. Motives

The discussion and series of experiments described in this work are motivated by a desire to challenge our perceptions about the roles of competition and cooperation in evolutionary processes.

In evolutionary theory, it is a given that if two entities cooperate in an altruistic manner in any way, it is because they themselves gain in the long run – whether personally, or through the genes of their offspring or kin. This is born of sound theory, but sometimes leads to a mind-set that sees competition and selfishness elevated to an unjustified status – as some kind of all-encompassing framework for life. This point of view can lead to over-generalised assumptions such as, 'all living things are in competition', and to an expectation that the ultimate product of evolution will be a singular and isolate elite. As this work will show, the competitive framework is too simplistic.

Is the emphasis on competition and the accompanying elitism necessary? Perhaps by taking a broader, more long-term view of evolution as a whole, rather than focussing on the short-term adaptation of a single species, we may enable a shift in perspective; where competition becomes a less significant aspect of the evolutionary model, and mutual benefit and cooperation more valuable tools for explanation. Such a shift may enable us to avoid misleading assumptionsabout the processes of Creation.

1.2. Comparing natural evolution and Genetic Algorithms

1.2.1. The standard Genetic Algorithm, and terminology

The basic principle of a Genetic Algorithm(GA) is to apply a selection pressure to a gradually varying population [Goldberg, DE, 1989].

The standard GA uses a string of characters (typically binary bits) to represent the genome of each individual in a population. Random genomes are generated for the initial population. The process progresses by repeatedly selecting individuals as 'parents' and producing 'offspring'. Various genetic operators (see below) may be employed to produce an offspring genome which contains more or less of the parental genome(s). The offspring will therefore, in general, inherit some of the characteristics of the parent(s) but will also introduce some variation. As in nature, selection pressure favours some variants over others; that is, variants that are "fitter" (see below)produce more offspring than other, less fit individuals. This is typically implemented by increasing the likelihood of fit individuals being selected for reproduction. The population size is usually kept constant by removing individuals at the same rate as new offspring are produced – this may be an entire population at a time, a 'generational GA', or on a one-for-one basis, a 'steady state GA'.

To introduce new genes, the genes of the parent usually have a chance of random mutation as they are passed to the offspring. In addition, various genetic operators can be applied to the genes during reproduction, producing transformations of the gene string. In sexual reproduction, cross-over allows a subset of the genes from two parents to be included in one offspring. This enables recombination of the genes that are present in the population, and is an important feature of the algorithm (see the explanation of 'building-blocks' below). The employment of the various operators implements 'genetic variance'.

'Fitness' is used to refer to an individual's reproductive advantage. In practice this is indexed by an 'evaluation function' or 'evaluation criteria' – some measure of the individual's closeness to a solution or success in performing a task. This is used to determine the probability of an individual being selected for reproduction. The relationship between this measure and the probability of reproducing determines the strength of the 'selection pressure' – at one extreme, a high selection pressure would see only the best (highest fitness) individual in a generation having any offspring. At the other extreme, with a low selection pressure, the best individual may only have a marginally increased probability of reproduction over the worst.

'Population stability' refers to the tendency of the population as a whole to vary gradually or quickly and is affected by the extent of genetic variance and the strength of the selection pressure.

A 'fitness landscape' is a concept used to visualise the relative fitnesses of genotypes. It is the surface described by the fitnesses of genotypes when arranged according to their genetic neighbours – that is, those that are close in terms of applications of the genetic operators. In the extreme, a rugged fitness landscape defeats the intent of the GA since the selective advantage of a parent bears no relation to that of the offspring, and the search is therefore random. In the ideal case, the fitness landscape describes a smooth dome, where the local slope is always a good indication of the direction to the global optima.

A crucial aspect of using GAs lies in designing an 'encoding schema' such that it provides a smooth fitness landscape. An encoding schema is a schema that determines how any phenotype will be encoded as a genotype. It therefore defines the relationship between genetic variance and phenotypic variance, and more significantly, the relationship between genetic variance and fitness variance and thus the ruggedness of the fitness landscape.

1.2.2. Compared with natural evolution

Genetic Algorithms are inspired by natural evolution. However, the standard GA, as described above, models only a narrow view of it: differential selection within a fixed niche – a competitive and convergent process. Natural evolution involves a continuously changing ecosystem, creating a dynamical system between the adaptation of one species and that of other species and/or its environment – quite distinct from a fixed niche. Natural evolution includes species which coexist and, moreover, species which are mutually beneficial – competition is not the only significant relationship. Furthermore, natural evolution includes increases in the number of species, evidenced by the heterogeneity of life on Earth. It is not simply a convergent process.

The appeal of GAs lies in their potential to provide complex engineering solutions with minimum human design effort. However, the success of a GA is crucially dependent on many factors: the encoding schema and evaluation function of individuals, and the balance between genetic variance, population stability and selection pressure.

No general methodology exists for establishing these parameters and thus GAs, far from being the 'something for nothing' hoped for, are ultimately dependent on considerable design effort in each new application [Cliff & Mataric, 1996].

Poorly chosen evaluation functions will yield rugged fitness landscapes and will provide inefficient search and/or poor solutions. Still worse, other evaluation functions may yield misleading fitness landscapes where the local slope consistently leads the population away from the global optima to dead-end local optima.

One approach to avoiding this is to utilise a succession of evaluation criteria, of progressively increasing complexity (eg. [Yamauchi and Beer 1994], [Harvey, Husbands, and Cliff 1994], [Gruau 1996]). 'Leading' a population to a global maxima using this method puts the onus of design back on the researcher to some extent and loses the benefit of avoiding design assumptions. The most promising method of automating this incremental process is to include the (succession of) evaluation criteria in the evolutionary process. This is done by using one evolving species to provide the fitness measure of another, in a mutual way. The success of a species is measured according to its behaviour with respect to another evolving species – hence coevolution, (for example, [Hillis 1990], [Harvey, Husbands & Cliff 1994], [Sims 1995], [Batali & Kitcher 1995], and [Oliphant, M, 1995]).

However, the researcher is still left with the task of designing the encoding schema [Cliff & Mataric, 1996]. The required flexibility of an encoding schema, that is, the resolution of morphological and functional description, has to be balanced with the resolution of evaluations (or simulated trials) and the level of desired behaviour. And specifically, has to be placed at such a level that the complexity of the desired solution will be tractable.

Despite their advantages over many other algorithms, it is not at all clear that GAs and coevolution, as they stand, will overcome the combinatorics of evolving complex systems [Cliff & Mataric, 1996].

Since the problem of choosing evaluation criteria may be solved by including it in the evolutionary process (via coevolution), it seems reasonable to look for a method that might solve the encoding schema problem by including it in the evolutionary process too. Whereas coevolution creates a succession of evaluation criteria, we now require a succession of encoding schemata or building-blocks – some method that might enable the creation of new, higher-level building-blocks from old. Pondering this, and examining the differences between natural evolution and GAs listed above, led to an idea that may enable the tractable evolution of complex systems. Perhaps unsurprisingly, it transpires that natural evolution 'thought' of it first. The formation of mutually beneficial groups and their subsequent 'encapsulation', as described shortly, is a vital aspect of the natural evolutionary process which has been thus far neglected in artificial evolution. The modelling of this process may enable artificial evolution to create complex systems that are otherwise unevolvable.

1.3. Mutually beneficial groups in evolution: Symbiogenesis

The 'major transitions in evolution' involve a process whereby heterogeneous groups of mutually beneficial individuals, or 'entities', to use a more general term, somehow become a new, more complex individual/entity.

> *entities that were capable of independent replication before the transition can replicate only as part of a larger whole after it.* [Maynard-Smith & Szathmary, 1995, p.4].

Such reproductively dependent groups are effectively new individuals – or at least, new entities with "a degree of individuality", [Maynard-Smith & Szathmary, 1995, p.5]

The 'encapsulation' of a group into a new individual (see C below) can be seen as creating a new building-block, with which to assemble groups at the next level. Thus we have a succession of building-blocks of increasing complexity – in other words, the process by which encoding schemata are defined is included in the evolutionary process. The advantage that this enables is similar to that conferred by problem decomposition in traditional search algorithms, as described in section 6.

See Table 1 for a summary of symbiogenesis.

Section 2 gives examples and evidence for the significance of symbiogenesis and mutually beneficial groups in nature. Then, in the following sections, the enabling mechanisms and selection pressures for the process are discussed. Various computational models, Ma to Md, each investigating one aspect of each step of the process, are used to illustrate key concepts. Finally, a summary of the circumstances where mutually beneficial groups may influence evolutionary processes and the implications for Evolutionary Modelling techniques are outlined.

1.4. Overview of discussion and models

A: Competitive/cooperative relationships and resource sharing: It is generally assumed that any two species in an environment will be competing for one resource or another and that competition is the most significant relationship to model. However, the assumption of competition between species is sometimes ill-founded. Where there are independent resources, diversity may be favoured over convergence. And even a straight-forward model of resource sharing, Ma, yields interactions that may be completely counter-intuitive in this respect – demonstrating mutual dependence as well as competition.

Table 1: Summary of symbiogenesis and illustrative models.

Steps in Symbiogenesis	Models to investigate and illustrate an aspect of each step
A **Formation of groups** A heterogeneous ecosystem contains many interacting species. This may include a group of species with mutually beneficial influences.	Ma **Population dynamics** A novel arrangement of species relationships with indirect resource sharing shows mutual dependence as well as competition.
B **Group stability** The stability of the group is maintained temporarily, though susceptible to exploitative individuals.	Mb **Spatial Compartmentation** A re-implementation of Boerlijst & Hogeweg's cellular automata shows spontaneous formation of spiral waves that provide resistance to parasites.
C **Encapsulation** The group becomes reproductively dependent and compartmentalised, thus forming a new entity/ a new species. These more complex individuals each exhibit the behaviour formerly exhibited by the group.	Mc **The Group Effect** An original model shows how the presence of mutually beneficial groups in an ecosystem changes the fitness landscape so as to enable the evolution of new species that emulate the group. This 'Group Effect' occurs without exceptional mechanisms – it enables encapsulation via ordinary mutation and selection.
D **Successive levels of recombination** The individuals that formed the group were the 'building blocks' of the new individual. The new species duplicates and diverges, forming a new collection of heterogeneous species, and the process repeats as from step A, being re-applied at increasing levels of complexity.	Md **Multi-level encapsulation** An original model demonstrates the application of the group effect, Mc, over successive levels of complexity.

B: The stability of cooperative relationships and group selection: Cooperative relationships are highly susceptible to 'cheaters' or 'parasites' – those that would exploit the benefit of the relationship. Parasites generally destroy the stability of the cooperative group and make evolutionary selection of the group as a whole impossible. So, such 'group selection' can only take place under quite specific circumstances – where the group is somehow compartmentalised. Mb, a re-implementation of Boerlijst & Hogeweg's work [Boerlijst & Hogeweg, 1991], demonstrates the spontaneous formation of 'spiral waves' which effects a spatial compartmentation that offers resistance to parasites.

C: The formation of new species – The Group Effect: The stability afforded by the previous mechanisms enables a degree of group selection but the long term stability and subsequent adaptation of such groups would be much stronger if it could be more concretely compartmentalised – forming a new individual/species and thus becoming a new unit of selection. To assist discussion, any process that enables the formation of a new species that emulates the former group shall be termed 'encapsulation'. Examples of this encapsulation in the literature utilise quite exceptional mechanisms – the formation of cell membrane, or horizontal gene transfer (see 2.1.1). However, a novel simulation of group interaction, Mc, demonstrates how the presence of mutually beneficial groups in an ecosystem changes the shape of the fitness landscape so as to enable the evolution of new species that emulate the group without such exceptional mechanisms. This 'Group Effect'[2] enables encapsulation via ordinary mutation and selection.

D: Algorithmic complexity – building blocks and successive reapplication of the building-blocks principle: The Building-block hypothesis or schema hypothesis is the principle which provides GAs with an advantage over other stochastic methods, eg. a population of simulated annealers. The re-combination of partial solutions through cross-over enables an advantage analogous to that conferred by problem decomposition in traditional search. The 'Royal Road' functions devised by Forrest & Mitchell [1995], were intended to demonstrate this principal advantage of GAs. The fitness landscapes and the expected reduction in the search space of algorithms which exploit the building-blocks is described. However, the Royal Road experiments served merely to highlight a significant problem with standard GAs – 'hitch-hikers', genes introduced via cross-over from successful individuals which corrupt the building-blocks that lesser individuals could have contributed. The group effect offers an alternative mechanism for recombination of building-blocks which avoids this problem. A modified Royal Road function where the building-blocks are local optima, and where many successive levels of re-combination are rewarded, is described. A novel simulation using this landscape demonstrates the group effect being re-applied over successive levels of complexity, Md.

[2] Named by analogy with the Baldwin effect, described in section 5.1

2. Symbiogenesis in nature

2.1. “The Major Transitions in Evolution”

The book of this title [Maynard-Smith & Szathmary, 1995] details many significant stages in evolutionary history. The transitions examined are not selected for the development of phenotypic qualities, eg. flight or vision, but rather for the development of new methods by which information/complexity is replicated. The authors stress a particular feature which is common to many of these transitions:

> “*entities that were capable of independent replication before the transition can replicate only as part of a larger whole after it*” [Maynard-Smith & Szathmary, 1995, p.6].

This is the key step in symbiogenesis.

The authors give the examples shown in Table 2.

Table 2: Major transitions in evolution

The origin of	**Before transition**	**After transition**
Chromosomes	Independently replicating nucleic acid molecules	A set of linked molecules that must replicate together
Eukaryotes	The ancestors of mitochondria and chloroplasts were once free-living prokaryotes	They can now replicate only within a host cell
Sex	The first eukaryotes could, presumably, reproduce asexually and on their own	Today, most eukaryotes can replicate only as part of a sexual population
Multicellular organisms	The cells of animals, plants and fungi are descended from single-celled protists, each of which could survive on its own	Today they exist (outside the laboratory) only as parts of larger organisms
Social groups	Most insects can replicate individually/sexually	Individual ants, bees, wasps and termites can survive and transmit genes only as part of a social group: the same is effectively true of humans

adapted from [Maynard-Smith & Szathmary, 1995, p.6].

2.1.1. Evolution of eukaryotic cells

Lyn Margulis has proposed the now familiar theory of the evolution of eukaryotic cells from prokaryotic cells based on symbiosis:

> “*...different bacteria form consortia that, under ecological pressures, associate and undergo metabolic and genetic change such that their tightly integrated communities result in individuality at a more complex level of organisation.*” [Margulis, 1995], p.136.

In this case the encapsulation of the group into a new individual occurs via horizontal gene transfer [Bull & Fogarty, 1996].

2.2. Natural mechanisms for increasing complexity

Figure 1 shows two (of three) methods by which complexity, or more specifically, information content of the genome, may increase in natural evolution.

Figure 1: mechanisms for increasing complexity

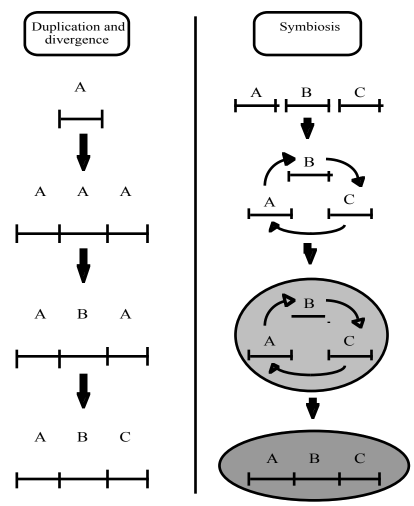

[adapted from [Maynard-Smith & Szathmary, 1995] p.11

The duplication and divergence scenario is the familiar one and is responsible for cell differentiation in multicellular organisms, for example. The symbiosis scenario is the subject under discussion.

2.3. Hypercycles

Eigen [1971] elucidated a paradox in evolutionary history: "...no enzymes without a large genome, and no large genome without enzymes" [Maynard-Smith & Szathmary, 1995].

The copying fidelity for non-enzymatic replication cannot be greater than 0.99, (99%) hence the selectively maintainable amount of information cannot be appreciably greater than 100 bits.

> the selectively maintainable amount of information N is given by,
>
> $N < \ln s/(1-q)$,
>
> where q is the copying fidelity and s is the 'selective superiority of the master' in the 'quasispecies' (from [Maynard-Smith & Szathmary, 1995] p.46)

"To increase this number, a replicase enzyme is needed. The smallest genome able to code for such an enzyme, and for the necessary translating machinery, would contain a number of genes whose total length would greatly exceed 100 nucleotides", ([Maynard-Smith & Szathmary, 1995] p.49). The solution to this paradox is found in *hypercycles* [Eigen & Schuster 1979]. These are *autocatalytic* sets [Kauffman 1986] of replicating molecules. That is, self-replicating molecules which catalyse (assist) the reproduction of other molecules in a mutual set.

Figure 2 shows a schematic of a five member hypercycle. A to E are self-replicating molecules and E catalyses the replication of D, which catalyses C, and so on, to A which completes the cycle by catalysing E.

Figure 2: five member hypercycle

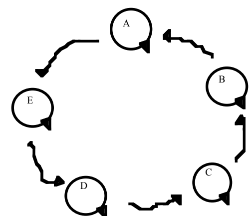

It is the hypercycle which enables the “stable coexistence of independent information carriers” which together code for enzymatic replication, and thus provides a solution to this ‘catch 22’.

In more general terms though, the hypercycle demonstrates a symbiotic relationship where reproductively independent entities become reproductively dependent. Mutually beneficial relationships are not exceptional. In fact, “ecosystems are full of [hypercycles]” ([Maynard-Smith & Szathmary, 1995] p.52) and they have, in at least this primitive case, enabled the evolution of complex entities that would otherwise been unevolvable.

The significance of such relationships to Artificial Evolution techniques may come only in similarly exceptional cases – but the following sections (particularly section 5) investigate the widespread influence of mutually beneficial relationships in evolutionary search.

3. Competitive/cooperative relationships and resource sharing

3.1. Competition and Cooperation

3.1.1. Coexistence – Diversity

Most coevolution simulations use two (or some other fixed number) of evolving species and enforce competition between them (eg. [Cliff and Miller 1995], [Sims 1995], [Hillis 1990], [Reynolds 1995]).

Natural evolutionary processes are much richer, of course. Most species inter-relationships are not directly competitive. In the scope of all species, direct competition, a predator/prey relationship, is incredibly rare – yet this seems to be the most common relationship modelled so far. It seems that the assumption of competition is quite widespread:

> *...the continuing popularity of the idea of "competition" in current ecological theory, about populations and communities, is a legacy from the neobiotic school that was prominent during the middle decades of the century and is still influential.*
>
> *The idea of competition, as it has been conceived by ecologists, is unrealistic when tested against what really happens in natural populations and communities.*
>
> *[Andrewartha and Birch, 1984], p131.*

The assumption that competition is the primary force in evolution is sometimes ill-founded. It is a product of modelling the dynamics of evolution on too small a scale – focussing on optimisation for a particular fixed niche, without recognising the dynamical nature of the relationship between the species and its environment.

Co-evolution simulations, as with standard GAs, are taken to be 'finished' when the populations are genetically converged and fitness stops increasing appreciably. Natural evolution has produced highly-heterogeneous ecosystems – it includes divergence as well as convergence. Enforced competition cannot support divergence.

Undeniably, optimisation of a species within a niche does occur, but this is meaningless without the context of the ecosystem. Competition of some kind is necessary to develop any form of efficiency within a species. Using enforced competition – where the fitness of one species is directly proportional to the failure of the other – a species that 'wanders off to do its own thing' has no fitness. Enforced competition can therefore not support divergence. A selection pressure is required that enables divergence and yet still optimises solutions on a local (niche) scale.

Resource-based competition

Divergence is a natural consequence of the fact that it may be easier to utilise a different resource than it is to compete for a single resource. Assuming finite resources, an exponentially increasing population will quickly overburden any given resource and this introduces competition between the individuals of the species. Adaptation that increases efficiency in utilising the resource (or destroying competition) will be selected for – hence, convergence. The alternative though, is adaptation to enable the utilisation of a different, as yet untapped, resource – hence, divergence. Thus, resource-based competition provides some balance between convergence and divergence (see [Pianka 1988]).

3.1.2. Cooperation/mutual benefit

Now we have seen that competition should not be presumed universally, and that divergence and coexistence are common, let us go a step further and examine the possibilities for cooperation or mutual benefit.[3]

An inter-relatedness of species within a divergent ecosystem may be maintained by virtue of the fact that resources are never truly independent, and so, it is generally assumed that all species must ultimately be in competition. However, the following model demonstrates that even where two species are ultimately dependent on a shared resource their populations are not necessarily in competition.

3.2. Ma Population dynamics:

A novel arrangement of species relationships with indirect resource sharing shows mutual dependence as well as competition.

In the case where one species preys on another the two species are engaged in direct competition. For example, vegetation that is eaten by caterpillars: the vegetation would survive better if it could avoid being eaten, and the caterpillars would survive better if they can eat more vegetation. (Notice that even this much is over simplified – perhaps the butterflies that result from the caterpillars are essential to fertilise the vegetation.)

Now suppose that there are two species which share a resource: for example, suppose rabbits and caterpillars both eat dandelions (and nothing else).[4] By virtue of the fact that dandelions are presumably a finite resource, and that the growth of either population is exponential, the rabbits and caterpillars are in competition. Thus far we are in agreement with our common assumptions about resource sharing and competition.

Now let us take the model one level further: suppose that there are two further species which prey upon the rabbits and caterpillars respectively (and exclusively), for example: foxes prey on the rabbits, and hedgehogs prey on the caterpillars, as in the Figure 3 below.

Figure 3: Five species with predatory relationships

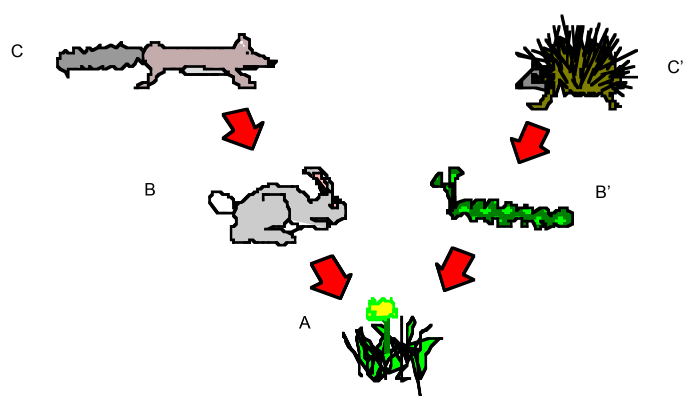

It may seem reasonable to extend the former argument and conclude that the foxes and hedgehogs are in competition by virtue of their (doubly-indirect) shared resource. It is this assumption, that all life is ultimately dependent on finite resources which are shared, that leads to the general conclusion that all species in an ecosystem are in competition. However, this assumption is naive.

Let us take an anecdotal examination of the population dynamics involved. Suppose the population of foxes increases, then the population of rabbits will decrease, which means there are less rabbits to eat dandelions and

---

[3] 'Cooperative' seems to imply intentional behaviour whereas 'mutually beneficial' does not imply intentionality, and for most of the discussion here 'mutually beneficial' is probably more relevant. However, the terms will be used interchangeably for the most part.

[4] Though it is only the relatedness of the species that is important here, and the biological correctness of specifics is not significant, the author apologises for any ignorance in this respect.

therefore more dandelions for the caterpillars... more caterpillars, more hedgehogs. So, an increase in foxes is correlated with an increase in hedgehogs – quite the contrary of the competitive relationship we may have assumed. On the other hand, we could argue that the population of foxes can only increase if the population of rabbits has increased, which leads to less dandelions for the caterpillars... less caterpillars, less hedgehogs. We see the weakness of anecdotal examinations and turn to a more concrete model. The following kinetic equations provide such a model.[5]

Density of species A is given by:

$$\frac{dA}{dt} = i - d_A A \tag{1}$$

where i is influx of As, d_A is decay rate of As.

This equation is then altered to take account of growth of Bs, when N As are consumed for the creation of every B:

$$\frac{dA}{dt} = i - d_A A - g_B BAN - g_{B'} B'AN' \tag{2}$$

where g_B is growth rate of Bs, B is density of Bs.

Density of species B or B¢ is given by:

$$\frac{dB}{dt} = g_B BA - d_B B \tag{3}$$

where d_B is decay rate of Bs.

Now, if we add Cs then density of B is altered to take account of the loss from the growth of Cs:

$$\frac{dB}{dt} = g_B BA - d_B B - g_C CBP \tag{4}$$

where g_C is growth rate of C, and P Bs are eaten per C.

Now, density of C or C¢ is given by:

$$\frac{dC}{dt} = g_C CB - d_C C \tag{5}$$

The above equations 2, 4 and 5 were simulated using finite element approximation (with complementary equations of 4 and 5 for the populations of B¢ and C¢, respectively).

Appendix B1 shows the equations in full.

3.2.1. Results

Various values for the growth and decay rates, and number_of_As_per_B were tried in these equations. At first these parameters were kept symmetrical for B:B¢ and C:C¢. It should be acknowledged that many combinations of values result in chaotic oscillations or extinction of one species or another. Saruwatari, Toquenaga and Hoshino [1995] demonstrate the instability of ecosystems with too high a level of interdependence.

---

[5] Equations 2 and 3 are adapted from Blomberg & Cronhjort (1994)

However, there were many combinations of values for these parameters that produced oscillatory populations which settled to constant population sizes. This is seen in the first part of graph 1a (and 1b).

To investigate the relationship between C and C¢, the growth and decay rates of C were altered after the populations had settled, and the effect on C¢ observed. If C and C¢ were in a competitive relation we would expect increases in the population of C to result in decreases in the population of C¢, and vice versa.

Graph 1a

In graph 1a (see appendix A1) the growth rate of C is decreased 3 fold after 1000 iterations. Counter intuitively, the graph shows the population of C *increasing* abruptly.

Let us return to the example species to make discussion easier. We have decreased the growth rate of foxes and the population has gone up. There is, in fact, a tiny decrease in fox population (about 0.02 or less on the scale used) but this immediately results in a large increase in rabbits, and subsequent increase in foxes. With time, the oscillations settle slightly and it is seen that the foxes remain better off. So, the growth of foxes had been keeping the rabbit population low. When this is relaxed both populations are better off. This in itself shows us that our intuitions about populations and their competition is not reliable.

As the populations settle again, it is seen that the competitive assumption between C and C¢ is correct in this case – as the population of C increases the population of C¢ decreases. And consistently, when the growth rate of C is subsequently changed to three times its original value, the population of C settles low whilst the population of C¢ goes high. C and C¢ populations are out of phase, and in competition (although the benefit/deficit for C is counter-intuitive, the relationship to C¢ is as expected).

Graph 1b

However, graph 1b shows a quite different relationship between the populations of C and C¢. In this experiment it is the decay rate of C which is altered rather than the growth rate. And in this case we can see that the population of C decreases as the decay rate is increased. (At least this much is as expect.) But now note that the population of C¢ decreasesalso.

As the foxes decrease, the population of rabbits explodes. The rabbits deplete the dandelions and thus detrimentally affect the caterpillars and subsequently the hedgehogs.

Here the populations of foxes and hedgehogs are in phase, not in competition. So, referring to our original anecdotal examinations, we see that the relationship between the success of foxes and the success/failure of hedgehogs depends on *how* the foxes succeeded – whether it is by decreased (*sic)* growth or by decreased decay. We see from the last experiment that if the hedgehogs had anything to do with it, although they would prefer that the foxes grow more quickly and live longer, quite the contrary of our intuitions.

3.3. Conclusions

It may not always be the case that resources are shared. Where resources are entirely independent there is much scope for non-competitive coexistence.

But, even when resources are shared the relationship between species is not a simple competitive one, and can in some cases exhibit instances of mutual benefit.

Limitations

The relationship between species sharing a resource is complex, not to say, chaotic. The experiments detailed here take a particular set of relationships and examine the results of a few changes in parameters. Another relationship that would be interesting to examine is circular ‘predation’, as occurs in hypercycles and in most, if not all, food ‘webs’. For other relationships and other parameter combinations, the mutualism displayed in the latter experiment might not hold or even be representative. Further work should investigate these alternatives, with the aim of clarifying which parameter sets result in dynamics that eventually stabilise, and which conditions will lead to populations being in phase or out of phase.

For now, we can be satisfied that competition should not be assumed, and that populations can be interrelated with mutual benefit.

4. The stability of cooperative relationships and group selection

4.1. Group selection

Even if, for whatever reason, there exists a group of individuals that confer mutual benefit on one another, it is not clear that the group as a whole can be selected for in evolutionary processes. The problem is that an individual that adapts to receive benefit from the group without giving benefit will be at an advantage. Such 'parasites' destroy the reciprocal arrangement and collapse the group. Thus a mutually beneficial group is unstable when members are under individual selection. In other words, since the 'unit of selection' is still the individual, the group is susceptible to exploitation by individuals. We cannot expect evolutionary processes to have the 'forethought' to recognise the long term benefit that an intact group confers – it is the reproductive advantage as seen from the perspective of the relevant unit of selection that drives the process.

In order to enable persistence and continued optimisation of the group, there must be some mechanism that isolates the influence of parasites. This can be achieved with a more general mechanism if the influence of all individuals is confined to small isolated groups or *compartments*.

Figure 4: group selection via dispersal and aggregation

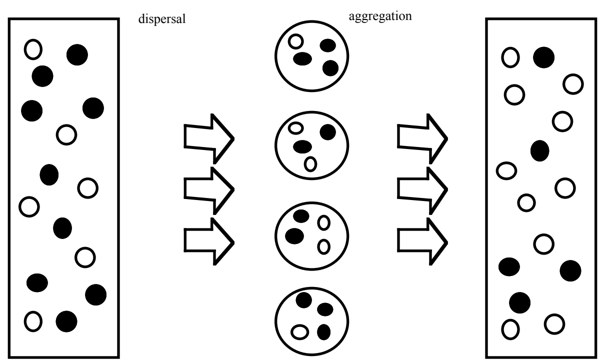

from [Maynard-Smith & Szathmary, 1995] p.55

In this model, a population contains two types of individual that are reproductively dependent, where one is beneficial to the other but this is not reciprocated – altruists and parasites. The population undergoes dispersal, stochastically forming small groups, with different proportions of the two types of individual. Those groups where the two types are present but the proportion of parasites is not too high reproduce most. Where one of the types is absent (ie. reproductive dependence is not satisfied) or where the ratio of parasites is too high there is less reproduction. So, although the ratio of parasites to altruists will increase in parasite-loaded groups, the *overall* proportion of altruists can increase.

An alternative model to isolate the influence of parasites relies on the spontaneous formation of spatial patterns. The formation of these patterns is demonstrated in Mb shortly.

The following chapter will discuss mechanisms whereby the temporary stability afforded by compartmentation or spatial patterns can be concretised by the creation of new individuals proper. If the group is encapsulated in a new individual (a permanent compartment) then it becomes a new unit of selection and can undergo subsequent adaptation whilst being robust to parasites.

4.2. Mb Spatial patterns to resist parasites

Re-implementation of Boerlijst & Hogeweg's cellular automata shows spontaneous formation of spiral waves that provide resistance to parasites.

A simulation of species distribution using cellular automata [Boerlijst & Hogeweg, 1991], shows the spontaneous formation of spiral waves. Such a spatial pattern restricts the influence of parasites and thus affords some stability for the altruistic group.

The model includes:

- spontaneous replication and decay,
- catalytic replication (for an autocatalytic set or hypercycle), and
- diffusion on a planer surface.

A cellular automata model uses a two-dimensional array of cells, each a finite state automata. Each cell's state is updated, either synchronously or asynchronously, according to the current state and those of neighbouring cells. In this case, the state of a cell can take one of seven values – empty or occupied by one of six species. If a cell is currently occupied it may become empty by spontaneous decay. If it is empty it may become occupied via the replication of a neighbour. The replication of a neighbour has a small chance of occurring spontaneously but a much bigger chance if neighbours adjacent to the direction of replication are catalysts for the replication of the speciesconcerned.

Figure 5: state transitions for cellular automata

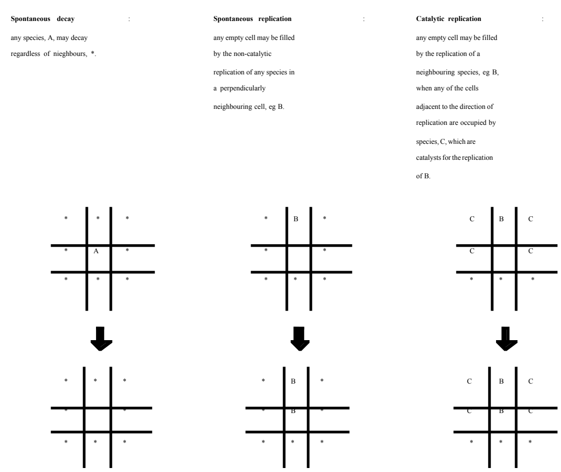

adapted from [Boerlijst & Hogeweg, 1991]

The six species that constitute the hypercycle are labelled 1 to 6. Species 1 catalyses the replication of species 2, 2 catalyses 3, and so on, with 6 catalysing 1. The replication of a species (into an empty cell) is probabilistic. The catalytic effect of having one catalyst in an appropriate cell contributes 10 times more than the spontaneous probability of a cell replicating.

For each perpendicular neighbour (north, south, east or west) the probability of their replication (into the empty cell) is the spontaneous probability plus the sum of the extra probability conferred by catalytic neighbours. So, in the catalytic reproduction figure shown above, B is a maximum 41 times (4 catalysts and spontaneous probability) more likely to replicate into the empty cell than it is in the spontaneous replication case.

Diffusion is implemented by taking a two by two block of cells and rotating them either clockwise or anti-clockwise at random. Adjacent blocks are treated in the same manner to cover the whole grid. In the next time cycle the blocks are scribed from areas which are a one cell diagonal displacement, thus implementing a shuffling of cells over larger areas in the long term.

Further implementational details are given in appendix B2.

4.2.1. Results

It is perhaps surprising that a global spatial structure should emerge from the above mechanisms. There is nothing in the algorithm that dictates any global trends; only local interactions are defined. However, graph 2a (see appendix A2) confirms the results of Boerlijst & Hogeweg [1991] and shows several spiral patterns which can be highlighted with the transparent overlay.

These spirals are the result of self-organisation. They result from the fact that replication is strongest on boundaries between areas of a species and its catalyst. Since the replication at such boundaries effectively moves the boundary the spirals rotate.

The patterns include double spirals which can be schematically represented as below.

Figure 6: schematic of double spiral

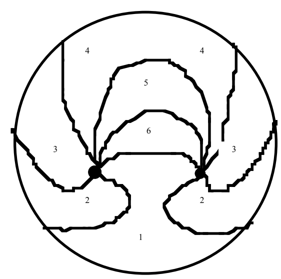

adapted from [Blomberg & Cronhjort 1994], p. 56.

Further results from Boerlijst & Hogeweg

Boerlijst & Hogeweg go further to demonstrate the robustness of these spirals by introducing parasites. They confirm that the spatial distribution and movement of the spiral waves restrict the effect of parasites. Parasites with the same reproductive advantages as existing species are introduced at random positions. They are catalytically supported by one of the normal species but do not offer catalytic support to any species. If they land in the arm of a rotating spiral their host species will pass over them in waves. But, since it takes many more than 10 time cycles for a complete rotation of the spiral, the parasite decays before the host returns.

Spirals can be destroyed by parasites if they land near the centre of the spiral but even in these cases the parasites infection is limited by the boundaries of surrounding spirals and does not spread further.

4.3. Conclusions

Mutually beneficial groups such as hypercycles are generally evolutionarily unstable. Since the group is not a unit of selection, it is the selective advantage of the individual that predominates. This makes such groups susceptible to exploitation by mutants that do not reciprocate the catalytic effect. Group selection can occur if the influence of parasites can be limited. Compartmentation or spatial self-organisation would enable this.

The above model demonstrates the spontaneous formation of spatial patterns. Such spatial patterns enable the influence of parasites to be isolated and thus affords mutually beneficial groups some resistance to them.

Limitations

It is not clear exactly what artefacts of the cellular automata model are required to enable the self-organisation demonstrated. Also, such organisation only applies to two-dimensional surfaces. This may be relevant to the interaction of primitive self-replicating molecules (on the surface of rocks) but this is a restricted application of mutually beneficial groups.

Compartmentation of some kind is required to enable even temporary robustness of the group. The next chapter investigates mechanisms whereby new species may be evolved – these being effectively a permanent compartment, and a new unit of selection – and can therefore undergo subsequent adaptation.[6]

[6] It should be noted that there may well be occasions when a distributed solution is better than a centralised one, eg [Husbands, 1992]. Thus the formation of groups without encapsulation may be an end in itself.

5. The formation of new species – The Group Effect

The stability afforded by the mechanisms discussed in the previous chapter enables a degree of group selection, but the long term stability and subsequent adaptation of such groups would be much stronger if it could be more concretely compartmentalised – encapsulated into a new individual/species, thus becoming a new unit of selection. This would then enable continued adaptation of the group.

Examples of this encapsulation in the literature utilise quite exceptional mechanisms. Perhaps the first form of permanent compartmentation would be the formation of cell membranes. Subsequently, the formation of eukaryotic cells, described by Margulis, which is the result of horizontal gene transfer (in otherwise asexual and distinct species of prokaryote cell; see 2.1.1). This rare phenomenon provides a method for the “consortia” to be permanently compartmentalised. But, it seems that such events are the result of quite exceptional mechanisms and circumstances.

However, the next section describes an original model, Mc, whereby the presence of mutually beneficial groups changes the shape of the fitness landscape to enable the evolution of complex entities that emulate the group without such exceptional mechanisms. This occurs without new membranes and without horizontal gene transfer – it occurs via ordinary mutation and selection of existing species, but these entities would not be evolvable without the presence of the group.

I have termed this process, where the presence of mutually beneficial groups *scaffolds* (see below) the evolution of new species which emulate the group, ‘The Group Effect’.[7] The group effect can be best described by drawing an analogy with the Baldwin effect [Baldwin, 1896].

5.1. The Baldwin effect

The Baldwin effect enables lifetime adaptations to influence the path of evolution without any Lamarckian inheritance of acquired characteristics. The ‘smoothing’ effect that lifetime learning can have on the fitness landscape has been demonstrated by Hinton and Nowlan, [1987] amongst others.[8]

Hinton and Nowlan [1987] use an “extreme and simple example” to demonstrate the Baldwin effect, as follows. Imagine organisms which may each exhibit a number of ‘phenes’.[9] Each phene has a corresponding phene with the opposite effect. Hinton and Nowlan suggest that the phene could correspond to connections in a neural network – a positive phene being a correct connection, and a negative phene an incorrect connection (or the absence of a correct one).

Let the number of phenes an individual can exhibit be P, and suppose that there is only one particular set of P phenes that confers added reproductive fitness. That is, in this set, all abilities are dependent on one another. And, suppose that possession of a phene is equally as likely as possession of a phene with the reverse effect: “in this worst case, there is no reasonable evolutionary path towards the [right combination of phenes] and a pure evolutionary search can only discover which of the potential [phenes] should be present by trying possibilities at random” [Hinton and Nowlan, 1987].

From the fitness landscape shown on the left of Figure 7, we can imagine a population wandering, unguided, on a flat plateau; where an individual that has, for example, even 19 out of 20 phenes correct, is not given any indication that it is close to the one phenotype that is of high fitness.

---

[7] This effect is previously described in the Adaptive Systems coursework for KBS, 1996. However, since the effect is now placed in the broader context of symbiogenesis, the description takes a significantly different emphasis.
[8] eg. [French and Messenger, 1995]
[9] The term used by French and Messinger [1995] to mean a part of the phenotype.

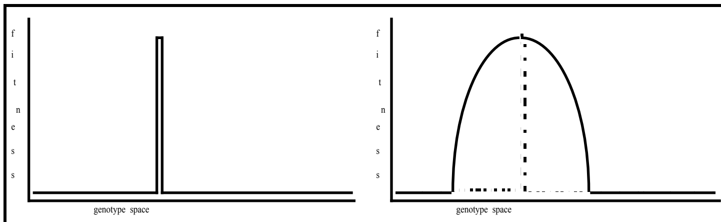

adapted from Hinton and Nowlan [1987].

Figure 7: 2D Interpretation of a fitness landscape

5.1.1. Introducinglifetimeplasticity

Now we introduce lifetime learning, or more generally, lifetime adaptation or plasticity, into the above scenario and examine how it affects the fitness landscape.

Suppose that individuals can have some phenes specified genetically and some phenes are left plastic, to be affected by lifetime adaptation.[10] There is a greater chance that an individual without any genetically specified phenes (ie. an entirely plastic individual) will find the correct set of phenes than that an individual with a genetically specified correct set will occur. Assuming that lifetime adaptation (phenotypic variation) is faster than genetic adaptation, this will be true even if lifetime adaptation is unguided (random).

5.1.2. Influence on fitness landscape

Now, if there is some reproductive cost to lifetime plasticity, or if it takes unproductive time to find the correct set of phenes during the lifetime, then individuals that have any correct phenes genetically specified will have increased fitness and will be selected for. The closer they are to a complete set of correct phenes which are genetically specified, the better. This is in contrast to the case before the introduction of lifetime adaptation where a genotype with nearly all the phenes correct was no better than any other unless it had the complete and exact set. Thus the lifetime adaptation provides a broader range of genotypes which are selected for and the shape of the fitness landscape can be pictured as in the right-hand part of Figure 7.

This new landscape enables the evolution of new species that contain the phenes found by lifetime adaptation. This occurs without a mechanism for the adapted phenotype to be inherited, which would be the biologically implausible ‘Lamarkianinheritance’.

5.2. Mc The Group Effect[7]

An original model shows how the presence of mutually beneficial groups in an ecosystem changes the fitness landscape so as to enable the evolution of new species that emulate the group. This ‘Group Effect’ occurs without exceptional mechanisms – it enables encapsulation via ordinary mutation and selection.

The group effect is now clearly demonstrated using an adaptation of the “extreme and simple example” used by Hinton and Nowlan [1987] to demonstrate the Baldwin effect.

---

[10]Plastic phenes are introduced as unspecified genes with probability 0.5, leaving 0.25 for positive and 0.25 for negative phenes.

5.2.1. Adapting the Baldwin effect example – Introducing group interaction instead of lifetime plasticity

Suppose that individuals can exhibit a variable number of phenes and that the phenes exhibited by an individual can be accumulated with those of another nearby individual or individuals. So, a group may be able to do some behaviour that a single individual cannot. In the extreme, a set of individuals with one correct phene each, may exhibit the complete set of correct phenes between them. As in the hypercycle, the effects of each individual 'chain together' with the next to form some beneficial process.

Suppose that the probability of exhibiting a phene is 0.5 and that for each phene an individual exhibits, possession of a phene with the reverse effect is equally as likely. This is equivalent to a genotype where genes may be either positive, negative or neutral/unspecified. So, the probabilities are 0.25, 0.25, and 0.5, for positive, negative and neutral phenes respectively.

To adapt the neural network model, it may help to imagine a network of connections which may be excitatory or inhibitive – then a positive phene corresponds to a connection of the correct type; a negative phene corresponds to a connection of the wrong type; and a neutral phene corresponds to no connection at all.

Without the influence of individuals in groups, the fitness landscape is the same as in the left of Figure 7. But, if we assume that the shuffling of individuals in the ecosystem enables the testing of more phene combinations than genetic variance does, then it is more likely that a complete set of phenes will be found via a complementary group than by an individual.

5.2.2. Influence on fitness landscape

Now, if there is some reproductive cost in forming large groups, or if it takes unproductive time to find the correct set of individuals, then individuals that have more correct phenes genetically specified will be selected for because they can form complementary groups with fewer members or form complementary groups more quickly/more of the time. As in the Baldwin effect, the closer they are to a complete set of correct phenes which are genetically specified, the better. And again, this is in contrast to the case before the introduction of group influences, where a genotype with nearly all the phenes correct was no better than any other unless it had the complete and exact set. Thus the group interaction provides a broader range of genotypes which are selected for and the shape of the fitness landscape can be pictured as in the right-hand part of Figure 7.

It is useful to think of the group as *scaffolding* the evolution of a new species. A species with many (but not all) phenes is selected for only when scaffolded by the presence of the complementary group members. As the number of genetically specified phenes increases, the group size can diminish whilst still providing a full set of phenes. Eventually, an individual with all phenes occurs and the group is no longer required.

The accumulative effect of the (non-neutral) phenes can thus enable a tractable search space for evolution. Drawing an analogy with the fact that the Baldwin effect does not require Lamarckian inheritance, the Group Effect does not require that genes responsible for the required phenes be inherited by one individual from another (ie. no horizontal gene-transfer). The reproductive advantage selects for individuals that contribute to the combined effect without communicating information about what additional phenes are required to the genotype. This means that the Group Effect could still be influencial in asexual organisms and even where groups are made up of different species.

Thus, the presence of mutually beneficial groups enables encapsulation in a broad range of cases and without exceptional mechanisms.

See appendix B3 for further implementational details.

5.2.3. Results

Appendix A3 shows the results of several runs of a simulation based on this scenario and compares evolutionary progress both with and without the group effect. The graphs monitor the relative frequency of positive, negative and neutral traits and show clearly that the group effect enables a marked increase in the frequency of positive phenes, and in some cases, the evolution of an individual with a complete set of positive phenes.

5.3. Conclusions

The formation of new species affords permanent stability to a group and enables continued adaptation and resistance to parasites. New species may occur via a number of mechanisms.

The Group Effect demonstrates that mutually beneficial groups can scaffold the evolution of complex individuals via ordinary mutation and selection. This shows that mutually beneficial groups may have a more widespread influence than the exceptional transitions described in the literature on symbiogenesis.

Limitations

It is important to recognise that neither the Baldwin effect nor the Group Effect demonstrate anything that would interest the engineer.[11] The Baldwin effect demonstrates the potential for one type of search to influence another – for lifetime learning to guide genetic assimilation (which is significant to the scientist). But, the computational effort required by the lifetime learning is no different from that which would have been incurred if the phenotype had been found by genetic search alone.

To quantify, Hinton and Nowlan demonstrate that the average number of positive phenes starts to increase after 30 or so generations – but with a population of a thousand and a thousand lifetime trials per individual that is 30.10^6 evaluations. For this 20 bit problem it would take 2^{20} evaluations to search the space exhaustively, which is 1.10^6 – much *faster* than the Baldwin effect. The fact that only 30,000 ‘genotypes’ were created is hardly interesting to the engineer when there were still 30,000,000 phenotypes to generate and test.

The group effect suffers from the same criticism. It takes 50 generations with a population of 1024 and an average of 930 trials per individual (making $47.6x10^6$ trials in total) to show an increase in positive phenes, which is somewhat worse than the Baldwin effect, and nearly 50 times worse than exhaustive search.

In the simulations thus far, the group effect offers no advantage in algorithmic complexity. So, although it is scientifically interesting to demonstrate that the presence of mutually beneficial groups in an ecosystem can influence the path of evolutionary search, we have not yet demonstrated any useful engineering techniques.

[11] Engineer is used to refer to a researcher motivated by the search for solutions to practical problems, as distinct from the scientist motivated by the search for understanding of nature.

6. Algorithmic complexity – building-blocks and the successive reapplication of the Group Effect

6.1. Rewarding partial solutions

In the problem described for Mc there is no explicit reward for partial solutions. A full solution must be found, albeit by a group of individuals rather than a single individual, before any fitness is apportioned to any individual. That is, the building-blocks for the complete solution are not discernible from the fitness function. The particular individuals that contribute to a complementary set of phenes are arbitrary, have no intrinsic fitness, and are found using no heuristic. This is just as in the Baldwin effect, which does not rely on anything other than random search in the lifetime learning – ie. the lifetime learning does not have any heuristics over and above those that are available to the evolutionary search.

So, finding the first successful group is completely random, just as finding the first individual that can learn successfully, is completely random. And therefore, the only reason the Group Effect (or the Baldwin effect), as described earlier scenario, has an advantage over the standard genetic algorithm is because the cost of evaluating groups (or evaluations during the lifetime) is not counted.

The Group Effect may be significant in these cases if generating and evaluating a group of simple individuals is 'cheaper' (or more likely) than generating and evaluating complex individuals. This may be the case in natural evolution, especially where an error threshold is involved (as in Eigen's paradox, section 2.3). This is the reason why changes in the method of information transmission are the focus of Maynard-Smith and Szathmary's work [1995].

The case where the fitness landscape is not a single pinnacle on a single plateau, but where partial solutions *are* rewarded, ie. the appropriate building-blocks are known, is quite different.

6.1.1. Building blocks

The building-block hypothesis or schema hypothesis is the principle which provides GAs with their advantage over other stochastic search methods. The advantage comes from trying combinations of entities already known to be useful. This combination of partial solutions enables an advantage analogous to that conferred by problem decomposition in traditional search. By finding the relevant building-blocks first, and then finding the right combination of building-blocks, the GA reduces the search space significantly. It is genetic cross-over that implements this recombination.

> To search for N blocks of k bits each, where any one complete block is rewarded, and blocks can be combined to find a complete set, then (as per the 'Idealised GA', [Holland et al 1995]) the expected time to find a complete set will be:
>
> $2^k \log_2 N$
>
> but, if there is no recombination, then the expected time to find a complete set will be:
>
> $2^k N \log_2 N$. (from [Holland et al 1995]).

Note that we are not comparing the expected time for 'search with reward for partial solutions' against 'search without reward for partial solutions', but rather, 'search with recombination of such partial solutions' against 'search without recombination.'

For this to be possible, partial solutions must be rewarded so that they can be selected for. So GAs are poorly suited for the fitness landscape with a single pinnacle on a single plateau. The overhead in maintaining a population of non-optimal solutions is likely to make GAs less efficient than simulated annealers if the recombination is not useful. And, with more complex algorithms that introduce additional overheads, such as those that utilise the Baldwin effect or the Group Effect, the situation is even worse, as detailed earlier.

We turn our attention then, to a different class of fitness landscape, one where partial solutions which combine as building-blocks for more complex solutions *are* rewarded. How likely is it that real problems are like this? It seems reasonable that by *assuming* that something that is good-so-far might be a building-block for something more useful, we stand a better chance of success than by random exploration. Or, to state it differently, it seems reasonable that the blind *combination* of entities will produce interesting results more often than the blind *modification* of them will.

6.1.2. Royal Road functions

The Royal Road functions, devised by Forrest & Mitchell [1995], were intended to provide fitness landscapes to exemplify this principal advantage of GAs.

Figure 8: Royal Road fitness function R1: from [Forrest & Mitchell 1995]

```
s1   =  11111111********************************************************;  c1    =8
s2   =  ********11111111************************************************;  c2    =8
s3   =  ****************11111111****************************************;  c3    =8
s4   =  ************************11111111********************************;  c4    =8
s5   =  ********************************11111111************************;  c5    =8
s6   =  ****************************************11111111****************;  c6    =8
s7   =  ************************************************11111111********;  c7    =8
s8   =  ********************************************************11111111;  c8    =8
sopt =  1111111111111111111111111111111111111111111111111111111111111111;  copt  =64

an optimal string is broken up into eight building-blocks. The fitness of an individual
genome is computed by summing the coefficients c1-8 of each schemata s1-8  that appears in the
genome.
For  example,  11111111000...=8,  1111111100000000111111100...=16.
```

The function described in Figure 8 describes a fitness landscape which is explicitly designed to reward building-blocks and thus lay out a 'Royal Road' – a path of ease – to exemplify the advantages of GAs. However, the design of a problem that exemplifies the advantages of GAs "turned out to be more difficult than we anticipated" ([Forrest & Mitchell 1995], p.4).

6.1.3. "Hitch-hikers"

The Royal Road experiments failed to provide a problem for which GAs performed better than hill-climbing! This is mostly because the building-blocks in R1 are not local optima but plateaux in the fitness landscape - and at least one type of hill-climbing still outperforms a GA on such a terraced landscape. A fitness landscape where the building-blocks are local optima would defeat hill-climbing. But, the Royal Road experiments also served to highlight a significant problem with standard GAs – 'hitch-hikers' [Holland et al 1995]. These are genes, introduced via cross-over from successful individuals, which corrupt the building-blocks that a lesser, receiving individual could have contributed.

Figure 9: Hitch-hiking corrupts building blocks

```
                                cross-over  point
                                        Y
A 11111111 11111111 00110100 11111111 11    |110001 00000001 00010000 - 3 blocks
B 00000000 11100111 11111111 00000001 10    |001000 00001010 11111111 - 2 blocks
_________________________________________   |_________________________________
C 11111111 11111111 00110100 11111111 11    |001000 00001010 11111111 - 4 blocks
```

Incorrect genes in A (shown in bold) 'hitch-hike' into C (carried by the momentum of As high fitness, if you like) and corrupt the first good block that was present in B, so the result has 4 blocks rather than 5.

The group effect offers an alternative mechanism for recombination of building-blocks which avoids the hitch-hikers problem. In the Group Effect, or any other method of encapsulation, the building-blocks are provided by individuals with a variable number of phenes. If there is some selection pressure favouring small individuals over large individuals with the same fitness then useless genes will be depleted. This means that when the genes of two individuals are combined there is less chance that destructive genes will corrupt existing building-blocks through hitching. A modification of Figure 9 will make this clearer.

Figure 10: How combination of individuals via Group effect avoids hitch-hiking

```
A 11111111 11111111 -------- 11111111 ----0--- ----1--- -------- – 3 blocks
B -------- ---1---- 11111111 -------- -------- --0----- 11111111 – 2 blocks
_____________________________________________________________________________
C 11111111 11111111 11111111 11111111 ----0--- --0-1--- 11111111 – 5 blocks
```

Although there is still a chance that incorrect genes in A or B will be encapsulated in C (in bold), useless genes or negative genes are generally absent (shown as -)[12] and therefore correct genes have a better probability of being encapsulated without corruption.

Figure 10 shows how combination of phenes through encapsulation is quite different from cross-over.

The following model uses this kind of recombination to demonstrate how recombination can occur through successive levels of complexity. An adapted and simplified Royal Road function is used to illustrate this process.

6.2. Md Multi-level encapsulation

An original model demonstrates the application of the group effect, Mc, over successive levels of complexity.

Two modifications to the Royal Road function provide a fitness landscape where there are many successive levels of re-combination rewarded and the building-blocks are local optima, and a third provides a simplified genome representation.

6.2.1. Modifying the Royal Road function

Successive levels of recombination

[Forrest & Mitchell 1995] detail a second version of Royal Road function, R2, that explicitly rewards successive levels of re-combination. In R2 the smallest rewarded block is 8 bits then 16, 32, 64.[13] In Md, only the largest block exhibited by an individual is rewarded – for example, an individual with a 32 block and an 8 block only receives benefit for the 32 block. This means that jumping to a higher fitness value is easier in the early stages than in the latter stages where recombination becomes increasingly more important (eg. to jump from 8 bits to 16 bits is easier than to jump from 16 to 32).

Providing local minima

The second adaptation of the fitness function is to add a cost for the size of the individual. That is, having 9 phenes is not only ‘no better than 8’ but is actually worse than 8. It seems reasonable that an individual that exhibits a useless trait should be less fit than one that does not. Figure 11 shows how the subtraction of the cost from the

[12] There is a tendency to assume that a gene must either be positive or negative, and cannot be unspecified/absent. This comes from the assumption that evolution is a process of function optimisation (with a fixed number of parameters) rather than open-ended search (with variable number of parameters). (see [Harvey 1994]).

[13] To be more consistent, Md, uses successive levels starting with 2 bits, then 4,8,16 etc.

reward function makes each building-block size into a local minima. This has the important effect that useless genes will tend to be absent so that encapsulation will be more successful than cross-over (see Figure 10).[14]

Figure 11: adapted fitness function

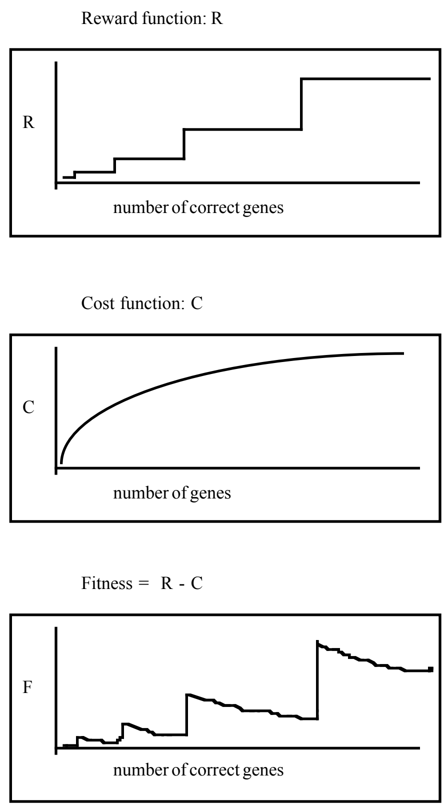

The reward function describes the benefit conferred by certain size thresholds.

The cost function describes the penalty for increasedsize.

The resultant fitness function shows how scaling successive levels becomes increasingly difficult.

Simplification

The third adaptation is one of simplification. Rather than only rewarding particular combinations of genes at specific loci, the adapted function rewards simply the number of correct genes. So, rather than a contiguous block of 8 bits being required, any 8 bits will be better than any 4. This means that we are simply rewarding certain thresholds in size. See Figure 11.

See appendix B4 for further implementational details.

6.2.2. Results

The results of using the group effect on this landscape are shown in the four graphs in appendix A4. The average fitness of the population is shown against time (in evaluations).

Graph 4a shows the progress of a GA without the Group Effect but with genetic recombination via cross-over. The first few jumps in fitness are made successfully (up to 64) but thereafter the population variance increases without finding the next level.

Graph 4b shows the progress of the population with the Group Effect enabling re-combination. As before the first few steps are scaled quickly. But, the fitness of individuals in the population continues to increase 'confidently' thereafter. Notice that there are several distinct steps in fitness increase that do not correspond to the reward function (eg. approximately, 46, 54, 75, 80). These correspond to the size of individuals that form successful groups (eg. 46*3 @ 128, 54*5 @ 256, 75*7 @ 512, etc.). It is these intermediates that enable the fitness of the population to continue climbing over the larger jumps to 128 bits. Graph 4c shows just the mean fitness of the population with Group Effect overlayed on the mean fitness of the population without Group Effect. Graph 4d compares the best in each population. Notice that the overhead of the Group Effect makes it worse than search without the Group Effect in some early stages but, in the latter stages, where the thresholds in fitness are further apart, it has a clear advantage.

6.3. Conclusions

The algorithmic advantage of GAs over other stochastic search methods comes from the re-combination of building-blocks in genetic cross-over. The Royal Road functions demonstrate that re-combination through cross-over can suffer from 'hitch-hikers' - mal-adapted genes carried on the momentum of generally fit genotypes which corrupt the building-blocks of less fit genotypes.

[14] It also means that the function is impossible for simple hill-climbing to solve (which is the shame of the Royal Road functions in the literature).

Re-combination via the Group Effect offers an alternative method of re-combination which avoids this problem.

A Royal Road function is adapted to include many successive levels of recombination, a cost for redundant genes and to provide a simplified representation of the genome. This is used to demonstrate the successive application of the group effect over successive levels of complexity and thereby complete the model of symbiogenesis summarised in Table 1.

For the engineer, it should be seen that the landscape used could not be outperformed by hill-climbing and that the simulation is a real comparison of GAs with and without the group effect. That is, unlike Mc, the evaluations for groups *are* included so as to make a proper comparison of algorithmic complexity.

Limitations

The fitness landscape used to demonstrate these principles is quite contrived. Building-blocks are explicitly rewarded and they are local optima. The encapsulation of individuals is greatly simplified by eliminating the specificity of genes and using simply the number of correct genes (or size of genome), see 6.2.1.

Thus the landscape is specifically designed to defeat hill-climbing and provide a landscape of re-combination that is simple for Group Effect to encapsulate.

Whether practical problems will have fitness landscapes with these features has not been demonstrated. However, some recent work by Gruau [1996], described following, provides an example of a problem where the standard GA was not suitable but which may have just such a landscape where symbiogenic mechanisms may be an advantage.

The task of evolving a controller for an eight-legged robot can be naturally decomposed to reward the behaviour of individual legs and then reward a suitable gate. However, application of the standard GA proved problematic. Consider a controller for a whole robot that has one leg behaving correctly - it is likely that the other legs are writing around randomly. In these circumstances, cross-over with an individual that has the correct behaviour for another leg is unlikely to produce an offspring with higher fitness than either parent because the re-combination will contribute more writhing than it contributes in well-behaved legs. That is, the genes ‘controlling’ the writhing legs would hitch-hike into the offspring and corrupt other successful building-blocks.

Gruau’s solution was to evolve the controller for a single leg in isolation and then manually replicate it before evolving the gate controller. However, the Group Effect, or some other method of encapsulation, would avoid the problems of the standard GA without manual intervention. The variable length genotypes would enable the evolution of individuals with one correct leg that make no specification - neither right nor wrong - about the other seven. It is easy to see that encapsulation of two complementary individuals of this type is more likely to produce successful offspring than the combination of two ‘writhers’.

7. Conclusions

7.1. How pervasive is the influence of mutually beneficial groups in nature?

Natural evolution is not simply a competitive and convergent process. Where resources are not shared there is plenty of scope for coexistence and as Ma shows, even when resources are shared, the relationship between species is not a simple competitive one, and can in some cases exhibit instances of mutual benefit.

There are several methods for enabling stability in mutually beneficial groups and offering some resistance to parasites. Mb shows one such method. In natural evolution encapsulation has involved the exceptional circumstances of horizontal gene-transfer or the formation of membranes. However, the Group Effect, Mc, demonstrates that mutually beneficial groups can scaffold the evolution of complex individuals via ordinary mutation and selection. This shows that mutually beneficial groups may have a more widespread influence than the exceptional transitions described in the literature on symbiogenesis.

In some sense the conditions outlined below for the application of symbiogenic mechanisms to engineering problems must also apply to natural evolution in some way such that the seemingly tortuous route of symbiogenesis will be favoured over more direct methods to increased complexity.

7.2. What is the applicability of symbiogenic mechanisms to practical engineering problems through Artificial Evolution?

The Group Effect when applied to landscapes like that described in Mc does offer an advantage in algorithmic complexity as it stands <u>if</u> there is some difference between the cost of generating and evaluating a group of simple individuals and the cost of generating and evaluating complex individuals, but otherwise it is a redundant overhead.

The case where partial solutions are discernible from the fitness function and building-blocks combine together is quite different. These are important characteristics for any kind of GA. However, the Royal Road experiments show that the standard GA suffers from the problem of hitch-hikers. The re-combination of building-blocks through the Group Effect offers a method to avoid this problem which occurs with re-combination in genetic cross-over.

A simplified problem with successive levels of re-combination and building-blocks which are local minima can be solved more efficiently with symbiogenic mechanisms such as the Group Effect than with standard GAs, Md.

In summary, symbiogenic mechanisms may provide an advantage where:

- there are partial solutions to a problem that confer fitness and combine as building-blocks in assemblies with greater conferred fitness,
- building-blocks are local optima,
- a selection pressure can be provided to maintain some diversity in the population and find the building-blocks (eg resource-based competition),
- members of mutually beneficial groups can be rewarded whilst mutations enable the group effect to encapsulate the group into a single individual.

Under these conditions the Group Effect can be re-applied over successive levels of complexity and may enable significant algorithmic advantage.

7.3. Future research

This work is in the early stages of investigation. The models described have each illustrated a key aspect of a step in the process of symbiogenesis but each has also had its limitations. The scope and applicability of each model should be fully investigated.

In particular, probably the most valuable direction for future research would be to identify a practical engineering problem and compare application of symbiogenic mechanisms against the performance of other algorithms. This would provide a concrete analysis of the advantages and show how the abstract concepts described can be applied.

In the meantime, researchers in Evolutionary Modelling and Artificial Evolution may at least give some consideration to using methods that:

- encourage diversity rather than using direct competition,
- enable group interaction and for evaluating the fitness of groups as well as individuals, and
- enable stability of mutually beneficial groups, be it spatial compartmentation or the formation of new species via the Group Effect.

The application of these methods in the cases described above may enable significant increases in the level of complexity that can be attained using Artificial Evolution.

8. Bibliography

[Andrewartha and Birch, 1984] "The Ecological Web", Andrewartha, H, Birch, L, 1984, University of Chicago.

[Baldwin, JM, 1896] "A new factor in evolution", American Naturalist, 30 441-451.

[Batali, J and Kitcher, P, 1995] "Evolutionary Dynamics of Altruistic Behaviour in Optimal and Compulsory Versions of the Iterated Prisoner Dilemma", Artificial Life IV.

[Belew, RK, 1989] "When both individuals and populations search" in Schaffer (ed) Proceedings of the third international conference on Genetic Algorithms. San Mateo, California, Morgan Kauffman.

[Blomberg, C and Cronhjort, M, 1994], "Modelling errors and Parasites in the evolution of primitive life: Possibilities of spatial self structuring" in *Cooperation and Conflict in General Evolutionary Processes*, Ed. Casti, JL and Karlqvist, A, John Wiley and Sons, NY.

[Boerlijst MC & Hogeweg P, 1991] "Spiral wave structure in PRe-Biotic evolution. Hypercycles stable against parasites", Physica D, 48.1, 17-28

[Bull, L & Fogarty, TC, 1996], "Horizontal Gene Transfer in Endosymbiosis", awaiting publication, (Uni. of West of England, Bristol).

[Cliff, D and Miller, GF 1994] "The coevolution of pursuit and evasion tactics in simulated robots", Cognitive Science Research Paper, University of Sussex.

[Dawkins, R 1976] "The Selfish Gene", Oxford University Press, NY.

[Eigen, M, 1971] "Self-organisation of matter and the evolution of biological macro-molecules", Naturwissenschaften, 58: 465-523.

[Eigen, M and Schuster, P, 1979] "The Hypercycle: A principle of Natural Self-Organisation", Springer, Berlin.

[Forrest, S and Mitchell, M, 1995] "Relative Building-block fitness and the Building-block Hypothesis", to appear in *Foundations of Genetic Algorithms 2,* Mogan Kaufmann, San Mateo, CA.

[French, RM and Messinger, A, 1995] "Genes, Phenes and the Baldwin Effect: Learning and evolution in a simulated population", Artificial Life IV.

[Gould SJ, 1985] "The flamingoes smile", WW Norton Company, NY.

[Goldberg, DE, 1989] "Genetic Algorithms in Search, Optimisation and Machine Learning", Reading Massachusetts, Addison-Wesley publishing Co.

[Gruau, F, 1996] "Evolution of neural networks to control locomotion of eight-legged robot"(?), awaiting publication.

[Harvey, I, 1992] "Species Adaptation Genetic Algorithms: A basis for a continuing SAGA", in *Towards a practice of autonomous systems, Proceeding of first European conference on artificial life*, Vareal, FJ, and Bourgine, P (eds), MIT Press/ Bradford Books.

[Harvey, I, 1993] "The puzzle of the persistent question marks: A case study of genetic drift", Cognitive Science Research Paper 278, University of Sussex.

[Harvey, I, Husbands, P, and Cliff, D, 1994] "Seeing The Light: Artificial Evolution, Real Vision.", in From Animals to Animats 3, editors Cliff, Husbands, Meyer and Wilson.

[Hillis, W, 1990] "Co-evolving parasites improve simulated evolution as an optimisation procedure" Physica D, 42:228-234,

[Hinton, GE and Nowlan, SJ, 1987]"How learning can guide evolution", Complex Systems 1, 495 - 502,.

[Holland, J et al 1995] "When will a Genetic Algorithm Outperform Hill-climbing?" to appear in *Advances in Neural Information Processing Systems 6,* Mogan Kaufmann, San Mateo, CA.

[Husbands, P 1992] "An ecosystems model for integrated production planning" CSRP 231

[Kauffman, SA, 1986] "Autocatalytic sets of proteins", Journal of theoretical biology 119: 1-24.

[Margulis, L, 1993] "Symbiosis in cell evolution", 2nd ed, WH Freeman and Co.

[Margulis, L, 1995] "Gaia Is a Tough Bitch" in "The Third Culture" Brockman, J, Simon and Schuster, New York.

[Mataric, M and Cliff, D, 1996], awaiting publication, (submitted to Simulation of Adaptive Behaviour, 1996).

[Mayley, G, 1996] "No Pain, No Gain: Landscapes, Learning Costs and Genetic Assimilation", submitted to EC and University of Sussex, CSRP 409(?)

[Merezhkovsky, KS, 1920] in Khakina LN (1992) "Concepts of Symbiogenesis: History of Symbiogenesis as an evolutionary mechanism", Yale Uni. Press.

[Oliphant, M, 1995] "Evolving Cooperation in the Non-Iterated Prisoner's Dilemma: The Importance of Spatial Organisation", Artificial Life IV.

[Pianka, ER, 1988]) "Evolutionary Ecology", 4th ed, Harper and Row, NY.

[Reynolds, CW, 1995] "Competition, coevolution and the game of tag", Artificial Life IV.

[Saruwatari, Toquenaga and Hoshino, 1995] "Adiversity - Stepping up trophic levels" in Artificial Life IV,.

[Sims, C, 1995] "Evolving 3D morphology and behaviour by competition" Artificial Life IV.

[Maynard-Smith, JM and Szathmary, E, 1995] "The Major Transition in Evolution", WH Freeman and Co.

[Waddington, CH, 1942] "Canalization of development and the inheritance of acquired characteristics", Nature 150, 563-565.

[Yamauchi, B and Beer, RD, 1994] "Integrating Reactive Sequential and Learning Behaviour Using Dynamical Neural Networks", in From Animals to Animats 3, editors Cliff, Husbands, Meyer and Wilson.