

How Symbiosis Can Guide Evolution

Richard A. Watson Jordan B. Pollack

Dynamical and Evolutionary Machine Organization
Volen Center for Complex Systems – Brandeis University – Waltham, MA – USA
richardw@cs.brandeis.edu

Abstract. Hinton and Nowlan have demonstrated a model of how lifetime plasticity can guide evolution. They show how acquired traits change the shape of the reward landscape in which subsequent genetic variation takes place, and in so doing encourage the discovery of equivalent heritable traits. This enables the seemingly Lamarckian inheritance of acquired characteristics without the direct transfer of information from the phenotype to the genotype. This paper draws direct inspiration from their work to illustrate a different phenomenon. We demonstrate how the formation of symbiotic relationships in an ecosystem can guide the course of subsequent genetic variation. This phenomenon can be described as two phases: First, symbiotic groups find solutions where individual organisms cannot, simply because lifetime interaction produces new combinations of abilities more rapidly than the relatively slow genetic variation of individuals. Second, these symbiotic groups subsequently change the shape of the reward landscape for evolution, providing a gradient that guides genetic variation to the same solution. Ultimately, an individual organism exhibits the capabilities formerly exhibited by the group. This process enables the combination of characteristics from organisms of distinct species without direct transfer of genetic information.

1 Introduction

Symbiosis, in its general definition, is simply the living together of different organisms. Often, in lay usage, the term is used to refer to the special case of *mutualism* where symbionts (organisms in symbiotic relationship) are mutually beneficial. Despite being undeniably common, the phenomenon of symbiosis, and especially mutualism, has for the most part been treated as a curio; a transient aberration on the otherwise relentless path of mutually-exclusive competition between species. In contrast, enlightened evolutionary theory recognises symbiosis as an integral process, and a fundamental source of innovation, in evolution. In its strongest form, symbiosis can lead to *symbiogenesis*: the genesis of new species via the genetic integration of symbionts [12], [9], [8], [10]. For example, eukaryote cells (from which all plants and animals are descended) have a symbiogenic origin [10].

The genetic integration of pre-adapted organisms is a fundamentally different source of innovation from the Darwinian accumulation of random variations. Computational abstractions of random variations under differential selection are well established in

evolutionary algorithms research. However, functional models for the interaction between the formation of symbiotic groups and the accumulation of genetic variation are under-researched. In this paper we make a modest start by modeling one mechanism by which symbiosis and genetic variation can interact. We do not pretend to model the biological details involved in any way, rather we provide an abstract computational model examining the dynamics of the adaptive mechanisms. And in this paper we commence with a model of an *indirect* mechanism whereby symbiosis can enable innovation, rather than the more radical process of symbiogenesis where symbionts are genetically integrated directly.

The mechanism behind our model is directly inspired by the work of Hinton and Nowlan in their 1987 paper [6], “How Learning Can Guide Evolution”. Their paper demonstrates the Baldwin effect [1]; a phenomenon whereby *acquired* characteristics can induce equivalent *heritable* characteristics. This seems like Lamarckian inheritance of acquired characteristics but it occurs without direct transfer of information from the phenotype to the genotype.

Hinton and Nowlan provide a simple and elegant abstract model that exemplifies how this process can occur and their model has been replicated and extended many times [2], [5], [11]. Here we have adapted their model by replacing learning with symbiosis; or more generally, replacing lifetime plasticity of an organism with lifetime interaction between organisms. Their experimental setup provides a convenient starting point and, moreover, an interesting comparison that assists us in understanding a more general concept that encompasses both phenomena.

Using this model we show that lifetime interaction can enable the evolution of organisms that would otherwise be unobtainable—or at least, would be very unlikely to occur. Our simulation of this phenomenon reveals two phases. Firstly, 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 combination of abilities via lifetime interaction of organisms samples a much larger set of variations than the relatively slow genetic variation from mutation. The first stage alone does not demonstrate the evolution of an organism that would otherwise not occur—rather we have simply selected a mutually beneficial group of organisms out of those that do occur.

In the second phase, after a group has found the solution and 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 will occasionally *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.

Thus, the abilities discovered by the symbiotic group become encapsulated in the heritable traits of a single individual. Yet this effect occurs without the exchange of genes—the symbionts may be distinct species. 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 the abilities are complete, the scaffolding is not required. This mechanism is illustrated in the following simulations.

The remaining sections of this paper are organised as follows. In Section 2 Hinton and Nowlan's learning scenario is adapted to our symbiosis model. Section 3 details the experimental setup and gives results. In section 4 we indicate a general principle that encompasses both the Baldwin effect and symbiotic scaffolding, and suggest implications for evolutionary computation methods. Section 5 concludes.

2 A scenario for symbiosis

Following Hinton and Nowlan's lead we describe an extreme and simple scenario where the combinatorics of the phenomenon 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. In Hinton and Nowlan's learning model the problem is to find the 20 correct connections for a neural network. For our symbiosis model we may imagine a chemical cycle with 20 steps. Each of the 20 steps must be catalysed by an organism correctly in order to get the chemical cycle going and to thereby confer reproductive fitness.

Where Hinton and Nowlan introduce lifetime plasticity to guide genetic search in this unforgiving landscape we shall use lifetime interaction with other organisms.

The interaction of organisms

Hinton and Nowlan suppose that some of the connections in the neural network will be left genetically unspecified and replaced with a switch that can make or break the connection during the lifetime of the organism. Here we suppose that an organism may have a neutral effect on a step in the chemical process, that is, the organism will neither prohibit nor catalyse the chemical step, and that this step may (or may not) be completed by some other organism in the environment. 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 traits of an organism in 20 genes where each gene has three alleles: *correct*, *incorrect* or *neutral* corresponding to catalytic, prohibitive (preventing completion of the cycle), or neutral interaction with a step in the chemical cycle. This unrealistic simplification enables us to see the mechanisms of interest more clearly but it is not integral to the result that follows.

Thus far we have followed Hinton and Nowlan's model exactly except with a different metaphor. Where they used correct connections, absent connections and plastic connections for a neural network we use catalytic, prohibitive and neutral

influence on steps of a chemical cycle. Now, where Hinton and Nowlan use lifetime learning to specify the plastic connections of the network we will substitute lifetime interaction between organisms to fill-in for the missing abilities of the cycle.

In the same way that Hinton and Nowlan use random search as their learning model because it makes the least assumptions necessary, we use the happenstance co-location of organisms to determine their symbiotic interaction since it makes the minimal assumptions. 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. Fig. 1 shows how the abilities of organisms are combined.

fifth organism:	--111----	0101--	0-0--
fourth organism:	-01---	0-0----	10-0-11
third organism:	1--10-11-	0----	10---1
second organism:	-1-0-1-001-	01-1--	0--
first organism:	00-0-11---	01-1--	0--1
combined abilities:	00100110010111100011		

Fig. 1. Combining the abilities of organisms. The 20 genes of each organism may take one of three alleles: correct, incorrect or neutral shown as 1, 0 and “-”, respectively. 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. Since every trait is specified by at least one of the first 4 organisms in this example, the fifth organism shown here is redundant.

Since the selection of, and ordering of, the organisms will be random, the details of this mechanism are largely inconsequential to the result that follows. One important feature, however, is that the fitness of the combined traits will be 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.

Evaluation

A key feature of the Hinton and Nowlan model is the fact that lifetime plasticity can search combinations of traits far more rapidly than genetic variation. They allow random search to test 1000 combinations of the plastic connections during the lifetime of one individual. Similarly, we allow the combination of abilities via lifetime interaction of organisms to be much more rapid than genetic variation. Accordingly we test 1000 random groups during the lifetime of an organism.

So, to evaluate an individual we divide its lifetime into 1000 time-steps. At each time-step a number of other organisms are selected at random from the ecosystem. The

number of individuals that are picked may be different for each time-step and the probable number is varied in the experiments that follow. The abilities of these organisms fill-in for the missing traits of the organism being evaluated, as described in Fig. 1. If the combined abilities of the group of organisms exhibits all required abilities correctly (and therefore enable the entire chemical cycle), then the organism receives a fitness increment, otherwise its fitness is unaffected. This is repeated for all 1000 time steps with a new randomly selected group each time. Overall, the fitness of an organism is given by $f=1+n$, where n is the number of time-steps when the organism in question forms a successful group with the organisms in its vicinity.¹

3 Experiments

The genetic model, the method of interaction, and the evaluation described above are iterated in a genetic algorithm (GA) [7]. Hinton and Nowlan choose 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 will find the solution but very likely that lifetime variation will. We continue to follow the experimental parameters of Hinton and Nowlan where applicable for the same reasons. We use a population of 1000 individuals, or in our symbiosis terminology, an ecosystem of 1000 organisms. Fitness-proportionate reproduction is applied generationally [7].

Since we are interested in the ability of symbiotic scaffolding to encapsulate the abilities of symbionts of different species, we will not use genetic recombination (crossover) in our main experiments. We shall return to this point below. Instead we will 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. Hence in generation 0 each organism will have about 50% neutral genes (as shown in Fig. 1).

Parts of the results that follow are largely predictable given the experimental setup—as in Hinton and Nowlan’s experiments the results are used to confirm our intuitions.

Experiment 1: A Crowded Ecosystem

In our first experiment we assume that the ecosystem is crowded. At any time-step there will be more than enough organisms in the immediate vicinity of an organism to fill-in all its neutral abilities. Naturally, when the organisms contain random genes they probably will not fill-in the steps in the cycle correctly. And, of course, if the organism

¹ In this detail we differ from the fitness function used by Hinton and Nowlan which is $1+19n/1000$, where n is the number of time steps *remaining after the first successful lifetime trial*. Our method gives the expected fitness of an individual to be directly proportional to the probability of success in a single trial whereas the Hinton and Nowlan model gives an expected fitness which is highly non-linear with respect to this probability [Harvey 1993].

being evaluated has any incorrect genes then it cannot form any successful groups. We build groups by accumulating organisms at random until all 20 abilities are specified (one way or the other). Since each organism specifies approximately half the full set of genes, the average number of organisms required to specify all traits is generally small.

Fig. 2 shows the number of each allele per organism averaged over all organisms in the ecosystem 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 20th generation a quite dramatic change takes place: the proportion of incorrect alleles falls close to zero whilst the number of correct alleles rises.² This is the point where symbiotic organisms become established and incorrect alleles are purged from the gene pool.

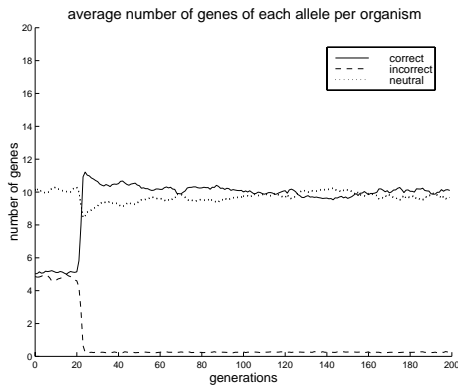


Fig. 2. Experiment 1. Number of genes of each allele per organism, averaged over all 1000 organisms at each generation. The first 200 generations are shown. Insignificant variation in the proportion of alleles occurs over subsequent generations (up to 1000).

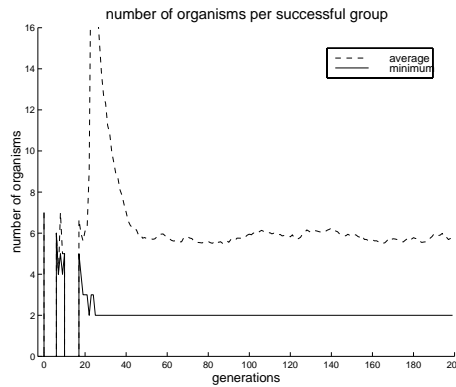


Fig. 3. Average and minimum group size of successful groups, experiment 1. Where no successful group is formed (as in some of the first few generations) both the average and minimum group sizes are shown as zero.

The effects of lifetime interaction shown in Fig. 2 have some qualitative similarity with the results of Hinton and Nowlan on the Baldwin effect. This is no coincidence; the substitution of lifetime interaction for lifetime plasticity is algorithmically insignificant in the initial stages where the abilities provided by other organisms are essentially random. However, once selection has taken hold and the population is predominantly devoid of incorrect alleles, the nature of the symbiotic variation is quite different from that of the random assignment used by Hinton and Nowlan. Specifically, the alleles supplied by symbionts are nearly all correct (whereas the traits supplied by random search remain in the original proportions). This means that the selection pressure for specifying increasing numbers of correct genes is very weak—in nearly all time-steps, other members of the group will correctly supply the required abilities. The

² Here, and more so in experiment two, the exact generation at which these sharp changes occur varies from run to run due to the stochastic nature of the experiment. However, the magnitude and general shape of the phenomena are reliable.

only significant pressure on an organism is that it should not specify any incorrect genes. A consequence of this complacency is that there is no pressure to get all correct abilities. Though Hinton and Nowlan's original work also shows no significant trend in decreasing neutral traits, replications of their work have shown subsequent increase in correct alleles over time-scales up to 500 generations [2], [5]. In this respect the influence of symbiotic groups is different from that of lifetime learning.

In Fig. 3 we see that the dramatic changes in Fig. 2 coincide with the establishment of groups that solve the complete cycle. We also see that the average size of successful groups increases dramatically. Given that organisms have approximately 50% of the correct alleles on average at this stage in the experiment, the fact that large groups are required to complete the set of abilities implies that many organisms have neutral alleles at the same loci. Informal observation also indicates that the population is somewhat converged at this time. By about the 50th generation the average group size falls to about six members which, given that the proportion of correct alleles has not changed significantly, shows that some complementary abilities have become established in the population. The smallest successful groups are of size two.

In summary, the crowded ecosystem in the first experiment shows the establishment of symbiotic organisms but does not demonstrate the entire scaffolding effect. We do not see the subsequent guidance of genetic variation to find an individual organism with the abilities formerly exhibited by the group.

Experiment 2: A Sparse Ecosystem

In our second experiment, we suppose the ecosystem is sparsely inhabited. Thus, the number of organisms in the immediate vicinity of the organism being tested is limited. Implementationally, we limit group sizes probabilistically where the limit is randomly selected from an exponential distribution. Specifically, the probability of there being exactly k members in a group is 2^{-k} , $k > 1$. 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. In short, an organism cannot rely on the availability of symbionts, and an organism that is more self-sufficient will receive a higher fitness on average.

In Fig. 4 we see the same dramatic trends seen around generation 20 in the first experiment, though here they occur later. But, we also see some quite different phenomenon thereafter. Unlike the first experiment we see a clear upward trend in the number of correct alleles in subsequent generations, and a significant increase starting at around 160 generations. Whereas the first sharp increase corresponds to the purging of incorrect alleles, the second corresponds to the purging of neutral alleles.

Examining the plots of successful-group size averages in Fig. 5 we see that, although there are a few instances of successful groups in the first 70 generations, the occurrence of successful groups takes longer to become established than in experiment one. This is reasonable, since in experiment two an organism has less interaction with other organisms and therefore fortuitous co-location is more sporadic. We also see that the average number of individuals per group does not escalate as acutely as observed under the evaluation scheme of the first experiment. Finally, we see that the rise in proportion

of correct alleles at 160 generations corresponds to the occurrence of individuals that are self-sufficient; i.e. a minimum group-size of 1 becomes established. In other independent runs there is considerable variation in the exact generation where symbiotic organisms become established and the generation which exhibits the first self-sufficient individual. Nevertheless, the overall effect is reliable.

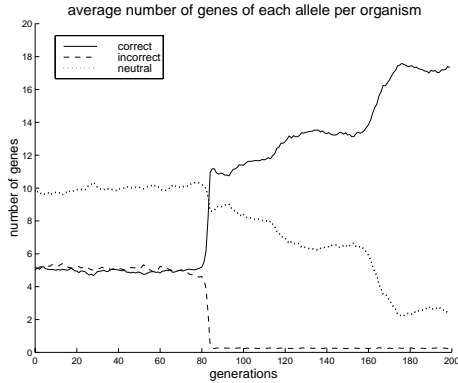


Fig. 4. Experiment 2. The average number of each allele per organism. All parameters as per the first experiment except that the size of groups is limited probabilistically.




Fig. 5. Number of individuals per successful group in experiment 2.

The difference between experiments one and two is that the high availability of symbionts in the first experiment produces complacency whereas the unreliable availability of symbionts in the second experiment provides a selection pressure favouring independent organisms. Thus in experiment two we witness the complete scaffolding effect: symbionts first enable the adaptive characteristics, then become obsolete. They have shaped the evolutionary search space so as to enable the evolution of organisms that perform the function formerly performed only by groups.³

4 Discussion

We can see in these results the emergence of a general principle that encompasses both symbiotic scaffolding and the Baldwin effect—specifically, a process of rapid variation guiding a process of slow variation. Either symbiosis or learning may guide subsequent genetic mutation. Both mechanisms are effective because their fast variation discovers a

³ Control experiments, not shown, included genetic recombination (one-point crossover) in addition to mutation. The results were essentially similar, though some effects were exaggerated probably due to stronger genetic drift [Harvey 1993]. It should also be noted that the recombination of genes via cross-over is quite distinct from the filling-in mechanism of lifetime interaction shown in Fig. 1; this distinction is to be expected in biological organisms also. Additionally, as expected, experiments run without lifetime interaction (with or without crossover) do not succeed in 1000 generations.

solution and then provides a gradient that guides genetic variation to the same solution.

From the biological point of view, the influence of symbiosis on evolution, whether indirect as in these models, or direct as in symbiogenesis, is of interest because it informs us of the origin of organisms we see in nature. Abstract computational models, such as that presented here and by Hinton and Nowlan, inform biology only inasmuch as they illustrate the possible space of dynamics and interactions. From the evolutionary computation point of view, the influence of symbiosis in guiding genetic variation is interesting only if it provides inspiration for more effective algorithms (here abstractions, rather than specific biological details, are more informative).

For example: Does the combination of a fast variation mechanism with a relatively slow variation mechanism provide a method that is more powerful than either mechanism alone? Why might the encapsulation of a group into a single individual be computationally important? What is the algorithmic difference between symbiotic combination and sexual recombination? We address these three questions in turn below: the suggestions here have yet to be investigated.

The coupling of fast and slow variation methods provides a balance between exploration and exploitation. A fast, non-permanent variation mechanism enables low-cost exploration (lookahead). Subsequent encapsulation via a method of slow variation, with high-commitment, enables stability from which further exploration may take place without disrupting solutions that have been proven.

The encapsulation of the group into a single organism enables the opportunity for the process to recurse with a larger unit of variation. This implies that the process may be applicable to hierarchical building-block problems [13] where search progresses from bit-combinations to schema-combinations as the building-block hypothesis suggests [7].

The difference between symbiotic combination and sexual recombination, or crossover, as used in existing GAs, is two-fold. First, symbiotic combination occurs between distinct organisms, whereas sexual recombination occurs between similar organisms, i.e. from the same gene-pool and necessarily highly-converged. This perhaps suggests the use of recombination operators that mate similar organisms frequently (as used in existing niching methods [3]) but also mate dissimilar organisms on rare occasions. Second, symbiotic recombination is additive whereas crossover is ‘either/or’—that is, the results of symbiogenesis have the sum of the genes from the donors, but offspring from sexual recombination have approximately half the genes from each parent. This has important implications for respecting the integrity of the schema represented by the parents. Such recombination is more akin to the Messy GA [4] or more generally the Incremental Commitment GA [14]. It can also be seen that the symbiotic filling-in of Fig. 1. corresponds to a ‘competitive template’ in the Messy GA.

5 Conclusions

We have seen two instances of an adaptive effect whereby non-genetic mechanisms guide the genetic make-up of organisms by shaping the evolutionary landscape. The Baldwin effect, modeled by Hinton and Nowlan, demonstrates learning as the first

example of such a mechanism. In this paper we have adapted Hinton and Nowlan's model to illustrate a second example. Specifically, our experiments demonstrate how symbiotic scaffolding can guide the genetic make-up of organisms and lead to the evolution of organisms that would otherwise not occur.

Acknowledgments

Thanks to the members of DEMO at Brandeis for assisting us in the research process, and also to the reviewers for their constructive suggestions.

References

1. Baldwin JM, 1896, "A New Factor in Evolution," *American Naturalist*, 30, 441-451.
2. Belew RK, 1989, "When Both Individuals and Populations Search," in Schaffer JD, ed., *ICGA3*, San Mateo California: Morgan Kaufmann.
3. Deb K & Goldberg DE, 1989, "An Investigation of Niche and Species Formation in Genetic Function Optimization," *ICGA3*, San Mateo, CA: Morgan Kaufmann.
4. Goldberg, DE, Korb, B, & Deb K, 1989 "Messy Genetic Algorithms: Motivation, Analysis and First Results," *Complex Systems* 1989, 3, 493-530.
5. 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.
6. Hinton GE, & Nowlan SJ, 1987, "How Learning Can Guide Evolution," *Complex Systems*, 1, 495-502.
7. Holland JH, 1975 *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: The University of Michigan Press.
8. 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.
9. 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. Peterograd (in Russian, see Khakhina 1992).
10. Margulis L, 1992 *Symbiosis in Cell Evolution: Microbial Communities in the Archean and Proterozoic Eons*. 2d ed., W.H. Freeman, New York.
11. 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).
12. 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).
13. Watson RA, Hornby GS, & Pollack JB, 1998, "Modeling Building-Block Interdependency," *PPSN V*, Eds. Eiben, Back, Schoenauer, Schwefel: Springer.
14. Watson, RA, & Pollack, JB, 1999, "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.