

LEGAL NOTICE

© ACM, 2012. This is the author version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the Genetic and Evolutionary Computation Conference, ISBN: 978-1-4503-1177-9, DOI: 10.1145/2330163.2330180

<http://dx.doi.org/10.1145/2330163.2330180>

Open-ended Coevolution and the Emergence of Complex Irreducible Functional Units in Iterated Number Sequence Games

Benjamin Inden
Artificial Intelligence Group
Bielefeld University
Universitätsstr. 25, 33615 Bielefeld, Germany
binden@techfak.uni-bielefeld.de

Categories and Subject Descriptors

J.3 [Computer Applications]: Life and Medical Sciences—
Biology and genetics

Keywords

coevolution, complexity, number sequence games

ABSTRACT

We present three related number sequence games as simple models of coevolution and demonstrate that they produce escalating arms races and irreducible functional units of unbounded size. We argue that our results imply that the models also show unbounded evolutionary activity according to a previous formal definition. Furthermore, we examine the robustness of the coevolutionary dynamics under different parameter regimes. We propose number sequence games as benchmarks for coevolutionary algorithms, and make some suggestions on adjusting task difficulty and choosing selection methods in coevolutionary algorithms.

1. INTRODUCTION

Open-ended evolution is an interesting concept both from the perspective of theoretical biology and from the perspective of evolutionary robotics, which aims to automatically construct robot morphologies and controllers in an incremental way. A precise measure for evolutionary activity that was introduced several years ago has been used to demonstrate that many of the early artificial life systems do not have the potential for open-ended evolution [1]. Although it was later shown that a particular artificial life system called “Geb” does show unbounded evolutionary dynamics according to that measure [4], many questions remain about what exactly causes open-ended evolution to occur. One assumption that seems to be plausible is that coevolution

can be open-ended provided the environment in which it occurs and the genetic system employed are sufficiently rich, whereas open-ended evolution seems to be hard to achieve using a static fitness function. However, a sufficiently simple model that can still produce open-ended evolution would certainly be helpful for more systematic investigations of the conditions necessary for open-ended evolution and sustained increases in complexity (complexification). The type of complexification that we are interested in is not just an increase in the number of components. We are interested in the number of components that, taken together, solve a particular problem arising as a consequence of the need to survive and reproduce, or in other words, that can be seen as performing a particular function.

Previous research on coevolutionary number games (starting with [16]) may provide some inspiration here. However, this research was targeted at other questions about coevolutionary dynamics, and the used encodings (single numbers or vectors of a few component numbers) do not provide the potential for complexification. Previous biological research on “Gene for Gene” coevolution models and various related extensions [14, 3] has focused on the dynamics of one gene or a fixed number of genes with several (typically 2) alleles each that can make a host resistant to infection, or a pathogen virulent. In these models, the more powerful alleles typically incur fitness costs, and the distributions of alleles over time in the two populations are studied. These models are not concerned with questions of complexification over time or the difficulties of finding powerful alleles in the first place, and “arms races” [6] arising in these models typically consist of cycles where the frequencies of a finite number of different strategies in the populations increase and decrease.

Another approach towards coevolution comes from the Webworld family of models [7]. In the basic Webworld model, a species is characterized by a fixed size list of features. It interacts with species that have other features as defined by a randomly initialized feature interaction matrix. The model consists of phases where species sizes are computed iteratively by means of differential equations until the food web becomes stable. In between these phases, new species are created by randomly changing the feature lists of existing species. This model has been extended such that the number of features of a species can vary within bounds, and a growth in complexity has been observed in the extended model, although the reasons for that growth were not entirely clear [11].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'12 July 7-11, 2012, Philadelphia, Pennsylvania, USA.
Copyright 2012 ACM 978-1-4503-1177-9/12/07 ...\$10.00.

The Foodchain model [9] is perhaps closest to what we are going to present here. It models symmetric competitive coevolution between individuals that are strings of letters with a fixed length. Some letters can be used for attack, others for defense, and the rest has no function. Matching is done between sequences of attack and defense letters present in the genome of two individuals to determine which one gains from the interaction. Point mutations and duplications are used in an evolutionary process that leads to complexification, i.e., a growth of functional subsequences in the genomes.

Paths of Darwinian evolution have previously been classified as serial direct evolution, parallel direct evolution, indirect evolution by redundancy removal, and indirect evolution by change of function (also known as *exaptation*) [15]. In response to claims by proponents of “Intelligent Design” that evolution cannot produce “irreducible complexity”, i.e., functional units of many components with the property that removal of any one part destroys the function [2], the authors of the above classification have argued that the two indirect evolution paths can produce irreducible functional units. Exaptation, just like open-ended evolution, seems more likely to occur in a coevolutionary scenario than with a static fitness function. In an attempt to demonstrate the emergence of complex function by evolution with the artificial life system AVIDA [10], the authors have manually provided intermediate rewards for organisms performing a number of simple logical operations to finally arrive at an organisms that can perform the logical operation “EQU” (= “A equals B”). Of course, hand-designing intermediate rewards becomes very difficult once more complex domains are considered. Besides, it raises the question whether the most interesting issue is actually where these intermediate rewards come from. Here we intend to go beyond that approach and demonstrate that intermediate rewards can be provided by using coevolution.

Three related simple models of coevolution are presented in this article. The first is a model of competitive coevolution (parasitism), the second of cooperative coevolution (mutualism), and the third of commensal coevolution (where one species gains from the interaction whereas the other is indifferent to it). We demonstrate that the models can produce escalating arms races and irreducible functional units of unbounded size under certain conditions. (Here we use the term “arms races” in a loose sense such that everything used to gain from an interaction is considered part of the arms, regardless of the nature of the interaction, whereas originally the term was used for competitive interactions.) Although we do not directly measure evolutionary activity as defined previously [1], we argue that our results imply that the models also show unbounded evolutionary activity according to that definition. Furthermore, we examine the robustness of the coevolutionary dynamics under different parameter regimes. We discuss possible implications of our results for coevolutionary algorithms.

2. METHODS

2.1 Number sequence games

There are two populations in our number sequence games. In the parasitic and commensal model, individuals from the first population termed symbionts gain from the interaction, while those from the second, the hosts, are influenced in a

negative way or not at all. In the mutualist model, individuals from both populations gain from the interactions provided they meet certain conditions.

The genotype — and the phenotype — of a member of either population is a sequence of numbers. What biological phenomena could be modeled by these number sequences? On a molecular level, the numbers could represent enzymes that need to be expressed in a particular temporal order to achieve a given function. On a behavioral level, they could represent behavioral primitives. Typical parasitic scenarios include host defense enzymes, and corresponding neutralizing parasite enzymes, or behavioral primitives in a pursuit and evasion scenario. Typical commensal scenarios include microbes that live on the waste products of other organisms, to the chemical composition of which they have to adapt. For typical mutualist scenarios, we may think of pollination, in which case individuals from the symbiont population correspond to pollinators, and those from the host population to plants.

In every generation, each organism in one population is tested against every organism in the other population. In the parasitic and commensal models, symbionts gain a fitness bonus for every host organism whose number sequence they completely match. This means that there must be pairwise matches between corresponding host and symbiont numbers. This pairwise matching is performed until either only the host sequence ends (the symbiont wins), only the symbiont sequence ends (the host wins), both sequences end (the symbiont wins), or a wrong number event occurs in the host sequence (the symbiont wins). A wrong number event means that the defense produced by the host is ineffective because it violates externally provided constraints. To use such constraints in the parasitic and commensal models has proved to be necessary because otherwise the search space of the symbionts (which have to exactly match the host numbers) is much more difficult than the search space of the hosts (which can choose any number), which lead to the hosts escaping from the symbionts early in evolution in preliminary experiments. We define a match between two numbers as equality here, but other relations (like complement) would be equivalent as long as they do not change the number of possible solutions. If the host wins, it gets the corresponding fitness bonus in the parasitic model. In the commensal model, the host always gets a fitness proportional to the length of the subsequence in its genome that fulfills the externally provided constraints, which effectively means that a simple fitness functions with intermediate rewards is provided to the hosts.

In the mutualist model, an organisms gains a fitness bonus if it matches the whole sequence of the organism from the other population. This means that if their sequences have equal length, both will obtain the fitness bonus, but if one has a longer sequence than the other, then only the one with the longer sequence will win (provided, of course, that it completely matches the shorter sequence). No external constraints are imposed in the mutualist model because it is by nature symmetric.

Each fitness bonus has a value of 1.0. To arrive at the final fitness of an organism, these bonuses are added, and then a fitness contribution due to genome length $f_l(o) = c_{lp} \exp(-0.1 \cdot [genome\ length])$ is added, where c_{lp} is a constant that determines the costs of adding and maintaining further genes.

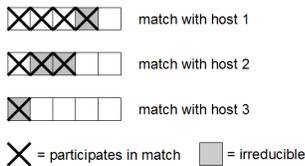


Figure 1: Measuring irreducible functional unit size (IFUS) in symbionts. In this example, the host population size is 3, and the symbiont sequence matches all host sequences. IFUS is defined as the maximum of the number of irreducible sites in a matching sequence, where irreducible means that there is no intermediate reward provided by other matches.

The nature of the constraint on the host number sequence in the parasitic and commensal model is as follows: The first number in the sequence must be 1, and a number at position $n + 1$ must have the value $v(n + 1) = (v(n) + 1) \% n_p$, where $0..n_p - 1$ is the range of possible gene values and '%' denotes the modulo operation. This means that only a fraction $1/n_p$ of the newly added genes will be effective. So the space of host solutions is constrained exactly as much as the space of symbiont solutions (where also only one number will match at a given position), and the only effective gene sequence of length $n < n_p$ takes the form $[1, \dots, n]$. It should be noted that neither the genetic system nor the fitness function (apart from this constraint) use any notion of neighborhood of numbers, so we could equivalently use any other constraint as long as effective solutions have the same probability in the search space under that constraint. Another important point is that, while we do provide a gradual path towards more complex solutions here, we provide it only to the host population. Furthermore, the host population in the parasitic model can only exploit the fitness gradient if the parasite population is at the right place in the sequence space to provide this information (whereas the host population in the commensal model can exploit information about the gradient at any time because its fitness is directly dependent on matching the external constraints). In section 3.3, we will introduce a more sophisticated method of setting constraints.

Simple fitness proportional selection (without any elite mechanism) is used as a starting point. Unlike many other selection methods, this one is known to be well-behaved with regards to game theory in the sense that all Nash equilibria of a given game are fixed-points of systems using fitness proportional selection, and all point attractors of systems using it are Nash equilibria of the game [8]. All experiments use a well-mixed population without spatial structure.

2.2 Genetic representation and operators

As already mentioned, the genome is basically a sequence of numbers. (For administrative and statistical purposes, we also provide each gene with a unique identification number as introduced in [13], but this has no influence on the model behavior as presented here.) Three mutation operators are used on these sequences: add a number (with probability 0.2), delete a number (with probability 0.1), and change a

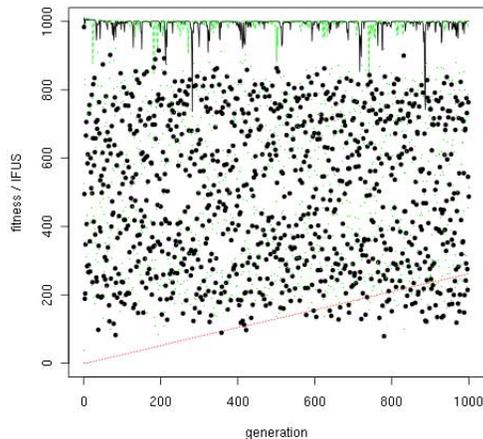


Figure 2: A typical run of the parasitic model. Maximal fitness of symbionts (solid black line) and hosts (dashed green line), mean fitness of symbionts (thick black dots) and hosts (thin green dots), maximal symbiont IFUS (red dotted line). IFUS increases linearly over time.

number (with probability 0.2). Values are always randomly drawn with uniform probabilities over the whole range. By default, the operations are only applied at the end of the sequence. Of course, the order of numbers in the genome could also be arbitrary if each gene also contained some information about the time of its expression. These ways of encoding the phenotype are equivalent.

All populations are seeded with a common ancestor that has a random sequence of length 1. Recombination is not used in the experiments reported here.

2.3 Measuring outcomes of coevolution

We are most interested in calculating *irreducible functional unit size* (IFUS) for organisms in the first population. This is done by iterating over all cases where an organism from the symbiont population matches an organism from the host population in a given generation (see Fig. 1). The sites in the symbiont genome that participated in that match are marked. Next, all sites that also participated in shorter matches in the given generation are unmarked. IFUS is then defined as the maximum of the number of marked sites over all matches. The highest IFUS found in the population is recorded in every generation. This in turn leads to two measures related to coevolutionary complexification: the highest IFUS ever reached during a run, and the number of generations before the maximal IFUS in the population finally dropped to zero. The latter value is equal to the total number of generations if no such drop occurs, or the population can recover from a temporary drop. This measure indicates for how long the symbiont population can compete before the host population completely escapes from it. We call this *effective symbiont competition time*, or ESCT.

3. RESULTS

3.1 Evolutionary complexification

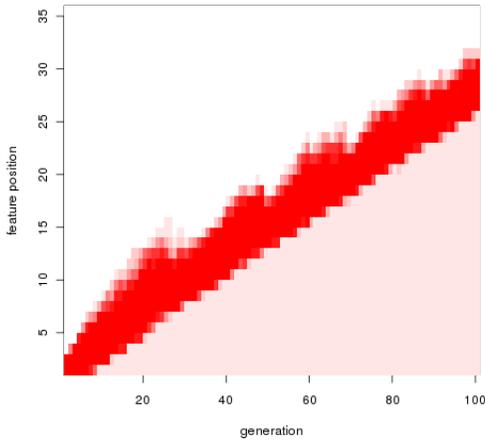


Figure 3: Sequence diversity in the host population for a typical run of the parasitic model. Diversity is color coded, with white corresponding to 0 and the most intense red corresponding to maximal diversity (in this case, 10) at a given position in the genome.

First we examine how the parasitic model behaves with different length punishment constants $c_{lp} = 0, 1, 10, 100, 1000$. As for all subsequent experiments, means of 20 runs with population sizes of 1000 and running for 1000 generations are reported. We set $n_p = 10$ for these experiments, so numbers can be in $\{0, 1, \dots, 9\}$. The maximal IFUS reached during coevolution using the different length punishment constants is 261.4, 261.1, 260.9, 260.0, and 23.9, respectively. ESCT is 1000 generations in all cases. This means that the parasite population can catch up with the host population in all cases, and irreducible functional units of similar sizes are produced as long as the costs of genes are not too high. The growth of IFUS is linear and therefore unbounded in typical runs (see Fig. 2). The mean fitness of the two populations oscillates with high frequency, while the maximal fitness values only occasionally drop much below the maximally obtainable fitness. So this is a fast arms race where one population is never much behind the other. Experiments not reported in detail here indicate that moderate changes in mutation rate change the speed of the arms race and therefore the maximally obtainable IFUS, but do not change the dynamics qualitatively. By looking at the host diversity (see Fig. 3), it can be found that only the last few positions in the sequence are variable, while the initial positions converge rapidly. The reason for convergence is, of course, that only a single value at any given position satisfies the external constraints.

When the mutualist model is run with $c_{lp} = 10$, an IFUS of 69.8 is reached on average. ESCT is 1000 generations in all cases, but the dynamics of IFUS is different in different runs: some experience intermediate breakdowns in complexity before recovery and renewed complexification (see Fig. 4). When the commensal model is run with $c_{lp} = 1$, an IFUS of 40.2 is reached on average. The symbiont population often cannot keep pace with the hosts, so IFUS breaks down. Sometimes there is an immediate recovery, but if not,

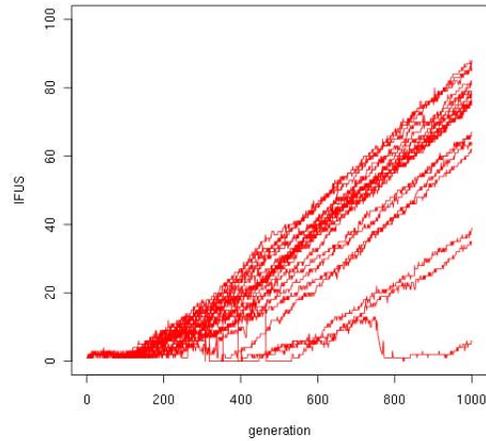


Figure 4: Complexification in 20 runs of the mutualist model. IFUS breaks down and builds up again in some runs. Otherwise, it grows linearly.

the IFUS stays at zero because the hosts have permanently escaped (Fig. 5), so ESCT is below 1000 in most cases, and IFUS does not grow unboundedly.

3.2 Robustness of complexification for different mutation and selection methods

Mutations have been implemented such that new numbers are always added at the end of a sequence. The justification for this is that the sequences represent not just the genotype, but features of the phenotype in this model, and it is certainly possible to conceive of genetic architectures where mutations are more or less strongly constrained to only add features without changing existing ones (on which they may depend), or situations where all other mutations are instantly removed by purifying selection. Nevertheless, one might wonder what happens in the more general case when numbers can be added anywhere in the sequence. In principle this should produce many individuals with a disrupted sequence and zero fitness, and make favorable mutations more rare. How does that affect coevolutionary dynamics?

In the parasitic model, complexification leads to average genome lengths of 97.8 in the final generation, and the symbionts remain competitive all the time. However, maximal IFUS remains close to zero. This can be explained by the effects of mutations: about 50% of the individuals are mutated, which means that about 500 mutations are introduced into the population in every generation. These are distributed randomly over the sequence length, so a host sequence in a given generation that satisfies the external constraints up to a position $p < length$ in the genome leads to many sequences in the next generation that satisfy the constraints up to positions $p_i \leq p$, thereby providing a path with many intermediate rewards towards the maximum for the symbiont population. This also results in a measurable increase in host diversity (see Fig. 6). Examination of the best individuals in the final generation reveals that evolution can maintain rather long correct initial sequences despite high mutation pressure. By contrast, the commensal

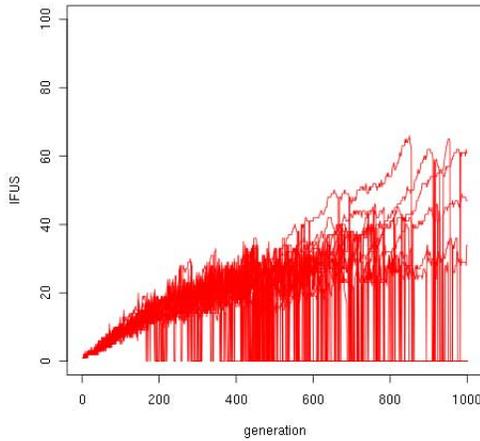


Figure 5: Complexification in 20 runs of the commensal model. IFUS frequently drops to zero, sometimes with immediate recovery, sometimes permanently.

model produces modest maximal IFUS (12.2), but permanent breakdowns occur in all runs, so complexification is bounded. The mutualist model reaches a maximal IFUS at a chance level of 3.2 only, and ESCT is usually below 1000. Complexification does no longer occur in this model.

Because of the above considerations that mutations anywhere in the sequence would produce many dysfunctional individuals, one might wonder whether a stronger type of selection perhaps could filter out these individuals and make all of the three models return to a complexifying dynamics. We investigated the effect of truncation selection (selecting a fraction of the population with the best fitness values, and assigning an equal number of offspring to each of them) with a truncation threshold of 0.01, which corresponds to selecting just 10 individuals from a population of 1000. The dynamics in the parasitic model remains the same qualitatively.

However, compared with fitness proportional selection, truncation selection significantly ($p < 10^{-9}$, t-test) increases maximal IFUS to 45.6 on average in the commensal model, though complexification is still bounded (see Fig. 7). In the mutualist model, truncation selection leads to complexification with a maximal IFUS of 90.8 on average. Like before, complexification seems to be unbounded in principle, but there are occasional breakdowns of complexity (see Fig. 8).

3.3 Effects of different constraints on the coevolving populations

So far, both populations have had search spaces with the same difficulty, i.e., the ratio of viable choices to possible choices of numbers has been the same for both populations. Instead, we can also choose two different numbers n_{c1} and n_{c2} and set the range of possible numbers to $[1, n_{c1} \cdot n_{c2}]$. Previously, only a single number as specified by an external constraint was valid at each position in the host genome. For the experiments in this section, the host number $n \in [1, n_{c1} \cdot n_{c2}]$ is divided by n_{c2} , and the result must then be a single number. In other words, the host number must be

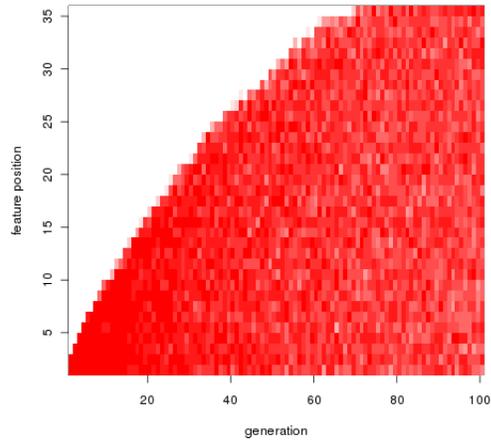


Figure 6: Host diversity for the parasitic model when mutations are allowed anywhere in the sequence. Compare with Fig. 3.

in the correct bin out of n_{c1} bins. Similarly, the symbiont and host numbers are divided by n_{c1} each, and the results must be equal for a match. In other words, we require that host and symbiont number are in the same out of n_{c2} bins (see Fig. 9).

We examined complexification for several different constraint ratios using the parasitic model with $c_{tp} = 10$, truncation selection with a threshold of 0.01, and two populations of 1000 individuals each running for 1000 generations (see Table 1). It can be seen that as the task of the symbionts becomes easier relative to the task of the host, complexification speed increases. However, complexification is fastest when the constraint ratios are exactly equal. This is remarkable given that even constraint ratios very close to equality do not gain an advantage. Why this is so remains unclear. Nevertheless, there is considerable robustness with respect to the constraint ratio. However, breakdowns in complexity often occur (see Fig. 10). When using fitness proportional selection instead of truncation selection, a bin size ratio of 100:100 yields maximal IFUS of 140.1 and no complexification breakdowns in most of the runs, whereas only very modest amounts of IFUS are produced with bin size ratios of 100:101 and 101:100 (16.2 and 18.3, respectively) and complexification breaks down early in these configurations.

4. DISCUSSION

4.1 Irreducibility of functional units

These results demonstrate that irreducible functional units can in principle arise by coevolution — even if the individual elements have costs. The matching sequences are irreducible because deleting or changing a single number from one of these sequences in evolved individuals would make the whole sequence dysfunctional with regards to their environment. Furthermore, these functional units are often so complex that their emergence in direct evolutionary scenarios without intermediate rewards is practically impossible. For example, the probability of a sequence of numbers between 0 and 9 and of length 260 that evolves in a typical

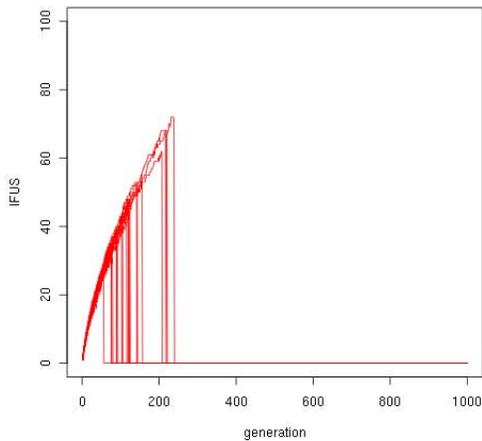


Figure 7: Complexification in 20 runs of the commensal model with mutations allowed anywhere in the sequence and truncation selection. IFUS eventually drops to zero in all runs because the symbionts cannot keep pace with the hosts.

parasitic model to arise by pure chance is 10^{-260} . By contrast, a single run consists of only 10^6 evaluations here.

To make the competitive and mutualist task approximately equally difficult for the host and symbiont populations, it has been necessary to constrain the solution space of the host population. One might think that makes the tasks equivalent to a simple evolutionary setup where a specified sequence has to be evolved and intermediate rewards are provided. But this is not the case for two reasons. Firstly, the information about the target sequence is provided only to the host population, but the sequence evolves in the parasite population as well. Secondly, the fitness function for the parasite population at any given time does not provide intermediate rewards for most of the sequence, but only for the latest additions to the sequence that have arisen in the host population. There are intermediate rewards only if all individuals that ever existed are taken into account (unless, of course, mutation provides intermediate rewards in every generations, as was the case in one of the models studied in section 3.2).

Furthermore, while we have implemented a simple identity function for matching between a host and a parasite, any different function would lead to the same results as long as the constraint ratio remains unchanged.

However, one limitation of the model is that complexification can only occur if the individual steps (in this case, adding a single valid number to the end of the sequence) are not too difficult. As the results in section 3.3 show, the parasitic model still produces large irreducible functional units if the range of possible numbers is extended to 100 or more, but complexification becomes less stable, and if the search space is further enlarged, at some point the probability of finding the next correct number by a mutation is too small given the total number of evaluations. Where this point is will certainly depend on the population size among other factors. The abstract models introduced in this article can of course not replace studies that examine whether basic as-

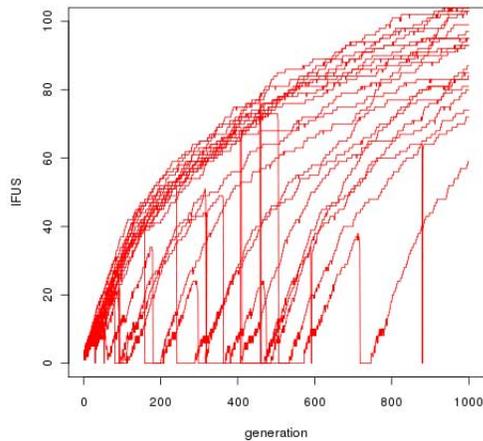


Figure 8: Complexification in 20 runs of the mutualist model with mutations allowed anywhere in the sequence and truncation selection. IFUS grows, although with occasional breakdowns.

bin size ratio	IFUS		bin size ratio
100:100	458.8		100:100
100:101	325.3	328.7	101:100
100:102	321.7	329.5	102:100
100:105	318.3	323.9	105:100
100:110	310.2	332.5	110:100
100:120	307.3	341.4	120:100
100:150	281.5	356.4	150:100
100:200	250.3	372.6	200:100

Table 1: Maximal IFUS reached for different constraint ratios in the parasitic model. Bin size ratios are given as (symbiont bin size ratio) : (host size ratio). A larger bin size means less bins and therefore makes the task easier for the respective population.

sumptions of the models regarding the nature of the search space are met in particular biological systems. However, they extend previous theoretical models of coevolution such that studies on complexification and irreducible functional units become possible, and the time required for finding feasible genes can be taken into account.

4.2 Open-ended evolution

According to [1], evolutionary activity (EA) can be measured as follows: all components of the system (e.g., genes) have an activity counter that is increased every time the component is used. The diversity is then the number of components with non-zero activity. The total cumulative evolutionary activity (TCEA) is defined as the sum of all activity counts, and the mean cumulative evolutionary activity (MCEA) is defined as TCEA divided by diversity. In addition, two levels of activity are defined such that the first level is close to chance activity level and the second is slightly above it. All components with an activity in that range are new to the system and have adaptive value, so the sum of their activity divided by diversity is defined as new

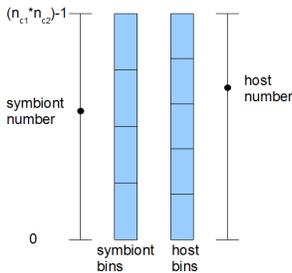


Figure 9: An example of how different constraint ratios can be achieved in number sequence tasks. A new number can have $n_{c1} \cdot n_{c2}$ different values. For the host population, it must be in the correct one out of $n_{c1} = 5$ bins. For the symbiont population, host and symbiont numbers must be in the same out of $n_{c2} = 4$ bins.

evolutionary activity (NEA). In order to determine chance levels for a given model, a *neutral shadow* model is created. It is like the original model with one exception: whenever selection is applied in the original model, a random choice is made in the neutral shadow. Using these statistics, a system can be classified as exhibiting unbounded evolutionary activity if its MCEA is unbounded and its NEA is positive (meaning that the limit of its integral over time, divided by time, is above zero). Modifications to these statistics proposed later [4] include resetting the neutral shadow periodically such that its population becomes identical to the population of the original model again, and using median cumulative evolutionary activity instead of the mean.

We have not measured evolutionary activity in our models, nor have we performed neutral shadow runs. Nevertheless, we argue that our models exhibit unbounded evolutionary activity if we define the components as genes and their usage as participation in a match. In the basic parasitic model (as in a number of other models and configurations), IFUS grows unboundedly. Because genes that participate in a match are considered active components with regards to EA statistics, that implies that the diversity of active components grows unboundedly. TCEA grows in an unbounded way, too (because old genes that participate in a match continue to be used), as does mean and median cumulative evolutionary activity (because most genes that participate in a match are old genes that continue to be used, so their activity counts keep increasing). Furthermore, NEA is positive because the chains are elongated all the time and the genes to be added at their ends will inevitably pass through the range of activity that has been determined for measuring NEA.

The neutral shadow run for our model can be performed as a thought experiment. In the neutral shadow, matches of longer sequences are extremely unlikely to occur because parasites with many matches have the same probability of survival as those with few matches. Therefore, component usage will be much lower. Furthermore, used components will be removed randomly over time due to deletions that destroy the match but have no effect in the neutral shadow. All in all, evolutionary activity in the neutral shadow will

be much lower. The same applies even if the shadow is periodically reset — degeneration of matching sequences will inevitably occur, and elongation of sequences will drop to chance level.

From the arguments in this section it can be concluded that our coevolutionary models are simple (to our knowledge, the simplest described so far) examples of systems with unbounded evolutionary activity. On the other hand, if we had a direct evolutionary setup where a single population was evolved to match an infinite sequence and intermediate rewards were provided, perhaps the same argument for unbounded evolutionary activity could be made. So perhaps unbounded evolutionary activity, while providing a quantitative measure for evolutionary processes, is not the only desirable quality of artificial evolutionary systems. Emergence of structures that are not completely reducible to environmental conditions is another desirable feature in many cases. At least our mutualist model does have that feature, too, because here the sequence does not depend on external constraints (the others can be argued to have it to some degree if, as in section 3.3, constraints are relaxed).

Unlike the parasitic and mutual model, the commensal model does not show potentially unbounded complexification. A plausible explanation is that a fitness gradient is provided all the time to the host population in this model, so the host population can continue complexification very fast even if there is a temporary breakdown in symbiont complexity, which leads to the symbionts permanently falling behind the hosts. This problem could be ameliorated, but not completely eliminated, by slowing down evolutionary speed in the host population.

4.3 Implications for research on coevolutionary algorithms

Research on evolutionary algorithms has greatly benefited from the widespread use of benchmark functions. We propose number sequence games as benchmarks for coevolutionary algorithms here. They can be used to obtain an empirical understanding of how different coevolutionary algorithms work, and how well they work in different kinds of situations. Once the properties of standard coevolutionary algorithms are better understood they in turn can be used to better understand the properties of different coevolutionary problems, and in particular of coevolutionary robotics problems, which are often too complex for direct analysis.

The results reported here provide some support for two design heuristics for coevolutionary algorithms: First, asymmetric arms races in coevolution work best if the tasks difficulty for the symbionts is not much larger than for the hosts. Therefore, it could sometimes pay to make the task more difficult for the hosts artificially to obtain a more balanced coevolutionary dynamics and open-ended evolution. Second, while fitness proportional selection will make the system reach game theoretic equilibria, other selection methods seem to achieve more robust complexification. The particular method used here, truncation selection, can be said to take intraspecific competition into account to get from the separately calculated fitness potentials to the final fitness values of individuals in the context of their population [12]. Truncation selection is known to efficiently reduce mutation load [5]. Truncation selection can be approximated by other models of selection, but may perhaps often be a plausible model of natural selection in its own right. A more sys-

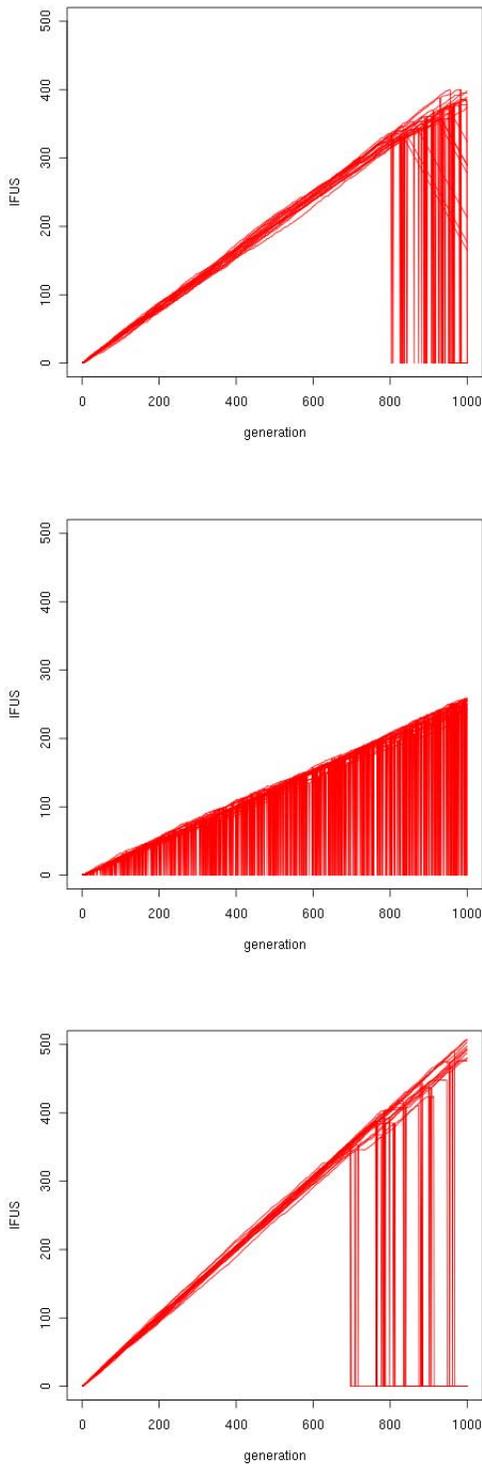


Figure 10: Complexification in 20 runs of the parasitic model with truncation selection and bin size ratios of (a) 200:100; (b) 100:200; (c) 100:100. Break-downs of IFUS occur most often when bin size ratio makes the task more difficult for the symbionts as in (b).

tematic investigation of the influence of different selection methods and their parameters on coevolutionary complexification is one area of possible future studies.

Acknowledgments

The author would like to thank Ipke Wachsmuth and Helge Ritter for hosting him in their departments, and acknowledges present funding by the Deutsche Forschungsgemeinschaft (DFG) in the Collaborative Research Center 673.

5. REFERENCES

- [1] M. A. Bedau, E. Snyder, and N. H. Packard. A classification of long-term evolutionary dynamics. In C. A. et al., editor, *Artificial Life VI*, 1998.
- [2] M. Behe. *Darwin's Black Box: The Biochemical Challenge to Evolution*. Free Press, 1996.
- [3] J. Bergelson, G. Dywer, and J. J. Emerson. Models and data on plant-enemy coevolution. *Annual Reviews of Genetics*, 35:469–499, 2001.
- [4] A. Channon. Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7:253–281, 2006.
- [5] J. F. Crow and M. Kimura. Efficiency of truncation selection. *Proceedings of the National Academy of Sciences*, 76:396–399, 1979.
- [6] R. Dawkins and J. R. Krebs. Arms races between and within species. *Proceedings of the Royal Society B*, 205:489–511, 1979.
- [7] B. Drossel, P. G. Higgs, and A. J. McKane. The influence of predator-prey population dynamics on the long-term evolution of food web structure. *Journal of Theoretical Biology*, 208:91–107, 2001.
- [8] S. G. Ficici, O. Melnik, and J. B. Pollack. A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *IEEE Transactions on Evolutionary Computation*, 9:580–602, 2005.
- [9] N. Guttenberg and N. Goldenfeld. Cascade of complexity in evolving predator-prey dynamics. *Physical Review Letters*, 100:058102, 2008.
- [10] R. E. Lenski, C. Ofria, R. T. Pennock, and C. Adami. The evolutionary origin of complex features. *Nature*, 423:139–144, 2003.
- [11] M. Maron and C. Fernando. Food webs and the evolution of organism complexity. In *Artificial Life X Workshop Proceedings*, 2006.
- [12] R. Milkman. A competitive selection model. *Genetics*, 74:727–732, 1973.
- [13] K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [14] J. N. Thompson and J. J. Burdon. Gene-for-gene coevolution between plants and parasites. *Nature*, 360:121–125, 1992.
- [15] R. H. Thornhill and D. W. Ussery. A classification of possible routes of darwinian evolution. *Journal of Theoretical Biology*, 203:111–116, 2003.
- [16] R. A. Watson and J. B. Pollack. Coevolutionary dynamics in a minimal substrate. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2005.