

Correlation Analysis of Coupled Fitness Landscapes

Wim Hordijk^{1,*} and Stuart A. Kauffman²

¹ Project Méthodes et Algorithmes pour la Bioinformatique, LIRMM, UMR CNRS 5506
- UMII, 161 rue Ada, 34392 Montpellier, France, wim@santafe.edu

² Director, Institute for Biocomplexity and Informatics, University of Calgary, Calgary,
Alberta, Canada, skauffman@ucalgary.ca

* Author for correspondence

Abstract

The correlation structure of fitness landscapes is a much used measure to characterize and classify various types of landscapes. However, analyzing the correlation structure of fitness landscapes has so far been restricted to static landscapes only. Here, we investigate the correlation structure of coupled, or dynamic, fitness landscapes. Using the NKC model of coevolution, we apply a correlation analysis on various instances of this model and present the results. One of the main goals of this paper is thus to show that a previously introduced correlation analysis can be successfully extended to coupled fitness landscapes. Furthermore, our analysis shows that this provides meaningful and interesting results that can contribute to a better understanding of coevolution in general.

1 Introduction

In 1932 Sewall Wright introduced the concept of fitness landscapes into evolutionary biology [1]. In a simplest version of this concept, one pictures a space of genotypes, with each genotype a single mutant distance away from its neighbors (i.e., neighboring genotypes differ at only one locus). This creates a high dimensional space of genotypes. To each point in this space a fitness value is assigned, assuming the limiting case that a species is isogenetic, occupying a single point in the space. This fitness value can be defined as one wishes, for example the propensity to have offspring, the ability to solve a certain given “problem”, or any other scalar measure. If the fitness distribution over the space of genotypes is considered, one obtains, in general, the picture of a more or less “mountainous” fitness landscape.

In a constant environment (static landscape), a species evolving by itself under mutation, recombination, and selection, explores this landscape in search of the highest peaks (or fitness values). For appropriate low rates of mutation and recombination, the population climbs to a nearby fitness peak and then remains trapped in the absence of outside perturbations.

One can measure a number of features of such more or less rugged landscapes, such as the number and distribution of local peaks, the distribution of their heights, the number of local peaks that can be accessed from a given initial condition, and other features. Important among such features is the correlation length of the landscape. Here one takes a random walk on the landscape, noting the fitness at each step, then computes the autocorrelation function of this series of fitness values, which gives the (average) correlation in fitness between two points a certain number of steps t apart. In general, smooth landscapes have long correlation lengths (there is still a significant amount of correlation for large values of t), random landscapes have zero correlation length, and rugged landscapes have correlation lengths that decrease as ruggedness increases. A complete correlation analysis of fitness landscapes, based on this random walk method, is briefly reviewed in section 3.

Contrary to the above (static) case, real evolution involves interactions between different species. Adaptive moves by members of one species deform the fitness landscape(s) of other species with which it is coevolving. Thus, a coevolutionary system is a general dynamical system without a potential function. The coevolutionary dynamics can exhibit the analogue of Nash equilibria where each species

is located on a peak consistent with the peaks of other species. Or oscillations may occur, where a move of one species causes the fitness peak of another species to move, so this second species has to move to this new location, causing the peak of the first species to move, so now this species has to move to this new location, causing the peak of the second species to move back to its former location, etc. The coevolutionary system may even exhibit chaotic behavior, when all species keep coevolving without settling down (the peaks are moving faster than the species can evolve).

So far, fitness landscapes of such coevolving systems, or coupled landscapes, have not been studied very much or very well. The purpose of this paper is to analyze the correlation structure of tunably rugged coupled fitness landscapes in an attempt to fill this lacune. In the next section, a general model for generating coupled landscapes (the NKC model) is briefly reviewed first. Section 3 then gives an short overview of a complete fitness landscape analysis that is extended and used here to study these coupled landscapes. In section 4 the general setup and model parameters we have used are summarized, while section 5 presents the results of our analysis. These results provide meaningful and interesting insights, summarized in the final section, that can contribute to a better understanding of coevolution in general.

2 Models of Fitness Landscapes

The tunably rugged landscape model we adopt is the NK model [2, 3, 4]. Here each gene in the genotype makes an independent fitness contribution f_i that depends upon the allele of that gene i itself and on the alleles of K other genes in the genotype. There are N total genes, and for simplicity it is assumed that the genotype is haploid and that there are only two possible alleles (0 and 1, i.e., the genotypes are bit strings). In the NK model, the K inputs to each gene (or *epistatic interactions*) can be chosen at random, or as the $K/2$ nearest genes on either side of the given gene (assuming the genotype is “circular”). This creates 2^{K+1} combinations of alleles for the $K+1$ genes that determine each gene’s fitness contribution f_i . These fitness contributions are simply assigned at random out of some statistical distribution as follows. For each gene i , each of the 2^{K+1} allele combinations is assigned a random value from the uniform distribution between 0.0 and 1.0. This creates a table of size $N \times 2^{K+1}$ with i.i.d. random values, from which the fitness contribution f_i of each gene can be determined for a given genotype. The overall fitness of this genotype is then simply the average over all fitness contributions, i.e., $f = 1/N \sum_{i=1}^N f_i$. For a detailed overview of the NK model, see the references given above.

The general properties of NK landscapes are as follows. When $K=0$, the landscape has a single peak, smooth sides and a long correlation length. When $K=N-1$, the landscape is completely random with nearly exponentially many peaks and zero correlation length. In between, the ruggedness of the landscapes increases and the correlation length decreases as K increases. Thus, by changing K (relative to N), we can tune the ruggedness of these fitness landscapes.

Coevolution among two species is now modeled by assigning to each gene in the genotype of species 1 “external” inputs from C genes in species 2, and vice versa, in addition to the K “internal” inputs. This raises the total number of allele combinations for each gene to 2^{K+C+1} . Again, fitness values are assigned to each combination at random from the uniform distribution between 0.0 and 1.0. The additional C inputs from the other species for each gene are also chosen at random. This coevolutionary model, known as the NKC model [5, 4], gives rise to tunably rugged coupled fitness landscapes. In general, if $C>0$, a move by one species on its landscape will deform the landscapes of its coevolutionary partners.

The next section will give an overview of a method for analyzing the correlation structure of (static) fitness landscapes, and how it can be extended to study coupled landscapes.

3 Fitness Landscape Analysis

In [6] a new statistical fitness landscape analysis method was introduced. This analysis is based on an earlier random walk method of [7], but additionally uses a time series analysis known as the Box-Jenkins approach [8, 9]. The advantage of this new method is that it provides a more detailed analysis of the landscape data, resulting in a complete stochastic model that expresses the correlation structure of the landscape, whereas previous methods generally only return one single number (being

some sort of average correlation or correlation length; see e.g. [7, 10, 11]) that summarizes the entire landscape structure. In [6] this landscape analysis was applied to the NK model, showing that the results have both explanatory and predictive value. In [12] it was subsequently applied to “less well-behaved” fitness landscapes, once more showing the usefulness and advantages of this more complete method over simpler methods. For a detailed introduction of the analysis we refer to the original paper. Here, we will give a brief overview of the main ingredients of the method, and describe how it can be extended to coupled fitness landscapes.

First, the analysis starts by generating a time series $f_t, t = 1, \dots, T$ of fitness values by performing a random walk on the fitness landscape. Starting with a random genotype, its fitness f_1 is calculated and recorded. Then a neighboring genotype in the landscape is chosen at random and its fitness f_2 is calculated and recorded, and so on up to some time step T . Note that the notion of a “neighboring genotype” depends on the move set that is considered (e.g., some sort of mutation or crossover). Here, we will only be concerned with point mutation (i.e., simply flipping one bit at a time).

Next, the Box-Jenkins approach [8, 9] is applied to this time series f_t . Briefly, this method first calculates the autocorrelations r_i and the partial autocorrelations a_i from the given time series. The autocorrelation r_i gives the (average) amount of correlation between two points in the time series that are i steps (called *time lags* in the language of time series analysis) apart. A value close to 1 means a high correlation, and a value close to 0 means hardly any correlation. The partial autocorrelation a_i can be interpreted as the average correlation between two points i time lags apart after the effect of the intermediate points on this correlation are taken out. Once the (partial) autocorrelations have been calculated, they are then used to identify and estimate an ARMA (autoregressive moving-average) model that adequately represents the stochastic process that generated the original time series. An ARMA(p, q) model is of the form

$$x_t = \alpha_1 x_{t-1} + \dots + \alpha_p x_{t-p} + \epsilon_t + \beta_1 \epsilon_{t-1} + \dots + \beta_q \epsilon_{t-q},$$

i.e., the current value in the time series depends on a weighted sum of p previous values plus a weighted sum of $q + 1$ white-noise terms. One could interpret p as the amount of “memory” that is in the system. Again, we refer to [6] for a detailed introduction and application of the method.

The above correlation analysis for fitness landscapes can be extended to coupled landscapes with only a slight modification. Suppose we have two coevolving species whose landscapes are coupled. First, a time series f_t of fitness values is generated as follows. Perform a random walk on the landscape of the first species, just as in the original analysis, and record the fitness value f_t at each time step. However, starting from a random genotype on the second landscape, move this second (or coevolving) genotype to a neighboring one in its landscape every m steps during the random walk on the first landscape. In other words, perform m steps of the random walk on the first landscape, recording the fitness values, then make a random move in the second landscape, perform m steps on the first landscape again, continuing to record the fitness values, make a random move in the second landscape again, and so on until T steps on the first landscape have been made. Once a time series of fitness values has been generated this way, the Box-Jenkins approach can be applied again, just as in the original landscape analysis.

This way a correlation analysis can be performed on coupled landscapes, introducing only one additional parameter, m , which can be interpreted as the relative rate of (co)evolution of the second species compared to that of the first species. If $m=0$, the second species does not evolve at all, and the analysis degenerates to that of a static (uncoupled) fitness landscape, which is useful for comparing the coevolutionary case against.

4 Experimental Setup

Using the NKC model as explained in section 2, and the extended fitness landscape analysis as explained in the previous section, we performed a correlation analysis of coupled fitness landscapes for different levels of couplings and rates of coevolution. This section provides details of the implementation and model parameters used, while the next section will present the results of the analysis.

For the implementation of the NKC model, we used a method of calculating fitness values similar to that described in section B2.7.2.4 of [13], and based on a previous implementation of the NK model by Terry Jones (personal communication). In this method, fitness values are not explicitly stored, but recalculated in an efficient way each time they are needed. This method has been tested extensively

in the previous implementation just mentioned. Furthermore, only nearest neighbor interactions are considered for the K (internal) epistatic interactions (where the genome is considered to be circular, so the first and last genes are each others direct neighbors). Finally, the C external epistatic interactions are chosen at random for each gene.

The parameter values we used for the NKC landscapes are as follows: $N=100$, $K=0,2,10,20,50$, and $C=0,2,10,20,50$. We did not consider values for K (or C) larger than 50, as previous results on regular NK landscapes have already shown that in these cases the landscapes become so rugged that there is hardly any correlation left, and they were not considered interesting for the present study. So, all in all we considered 25 different combinations of K and C, giving rise to as many different coupled landscapes.

For the random walks to generate time series of fitness values we used the following parameter values: $m=0,100,10,5,1$ and $T=100000$ ($=10^5$). We included the (degenerate) case of $m=0$ for comparison (and verification) with results on regular NK landscapes. So, on each of the 25 (coupled) landscapes we generated 5 random walks, each with a different value for the rate of coevolution m , and each walk consisting of 10^5 steps.

Finally, a simple C program was written to calculate the autocorrelations and partial autocorrelations of the time series of fitness values.

5 Results

In this section we present the main results of the correlation analysis as applied to the NKC model. If not explicitly mentioned, values for the model parameters as given in the previous section are assumed.

5.1 Comparison with regular NK landscapes

To verify the results on the NKC landscapes, they are first compared to those on regular (static) NK landscapes. Obviously, NKC landscapes with $C=0$ or with $m=0$ (regardless of the value of C) should generate the same results as regular NK landscapes with similar values for N and K. In the case of $C=0$, there is no coupling between the different species, so each species just evolves on its own landscape as if it was a regular, static landscape. When $m=0$, regardless of the value of C, there is a coupling between the different species, but the second species is not evolving (no updates or mutations). So, even though the second species does influence the fitness of the first species, this influence is not changing over time, and thus the landscape degenerates to a static one.

Figure 1 shows the autocorrelations r_i , as calculated from the time series of fitness values generated by the random walks, up to time lag $i = 50$ for regular NK landscapes for the various values of K. Clearly, the amount and length of the correlation decreases with increasing K. These results agree with results reported previously in [6], where it was also shown that in most cases the appropriate model to describe the correlation structure of these landscapes is an ARMA(1,0), or simply AR(1), model.

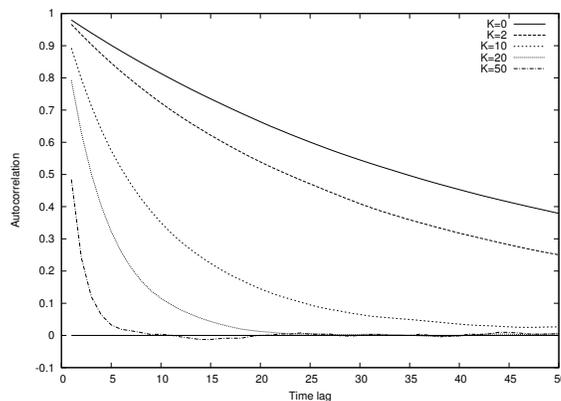


Figure 1: The autocorrelations r_i on regular NK landscapes for various values of K.

For comparison, figure 2 shows the same results, but this time together with those as calculated for the NKC landscapes for $C=0$ (left) and for $m=0$ (right). In both plots, the solid lines show the autocorrelations for the regular NK landscapes for the various values of K . In the plot on the left, the different dashed lines show the autocorrelations for the corresponding K values on the NKC landscapes with $C=0$ and different values of m . In the plot on the right, the dashed lines show the correlations for the corresponding K values for $m=0$ and different values of C . As the plots clearly show, they all agree very well with the regular NK landscapes, indicating that these cases indeed degenerate to a static landscape and thus form a base case for comparing coevolution with.

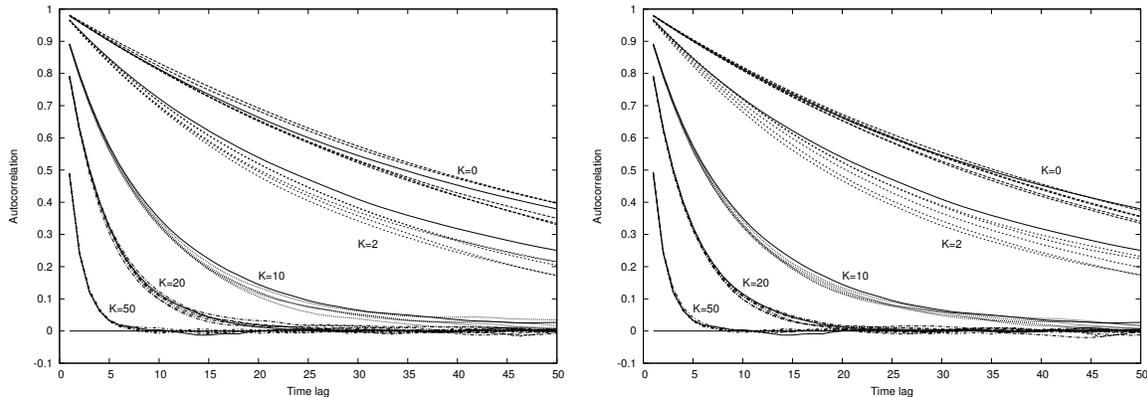


Figure 2: The measured autocorrelations on NKC landscapes for $C=0$ and different values of m (left, dashed lines) and for $m=0$ and different values of C (right, dashed lines). For comparison the correlations on regular NK landscapes for the various values of K are shown in solid lines in both plots.

5.2 Correlation on coupled landscapes

Next, we will look at the correlation structure of coupled landscapes. The analysis is split up in three cases: (1) a slow rate of coevolution (large values of m), (2) an intermediate rate of coevolution (intermediate values of m), and (3) a fast rate of coevolution (small values of m). The results are presented in the next three subsections.

5.2.1 Slow rate of coevolution

First, consider the case where the second species is evolving slowly compared to the first species, i.e., the value of m in our coevolution model is high. Figure 3 shows the main autocorrelation value (i.e., r_1 , or the correlation between neighboring points in the landscape) for the different combinations of K and C in a grey-scale plot for $m=100$.

As the plot shows, the value of C does not have any significant influence on the first autocorrelation. The color within one particular column (i.e., for one particular value of K) does not change for different values of C , and the main correlation value depends only on the value of K .

To have a more detailed view, figure 4 shows the autocorrelation plots for $m=100$ for $K=2$ (left) and $K=10$ (right) for the different values of C . Here, a subtle but surprising effect shows up for $K=2$. All the dashed lines (i.e., autocorrelations for $C>0$) are clearly above the solid line ($C=0$). It appears that for small values of K , a slow rate of coevolution actually increases the overall correlation length of the landscape as compared to the static case! To confirm this behavior, some additional experiments were performed with $m=50$ and $K=2$ (results not shown). The same situation occurs here as well (increase in correlation), except for the $C=50$ case, for which the autocorrelation curve is now below that of $C=0$. Thus, this increase in correlation seems to be a real phenomenon for low K and a slow rate of coevolution.

Indeed, it is known from several studies using genetic algorithms to solve some given problem, that introducing some form of coevolution can actually lead to a slight increase in the quality of the obtained solutions (see e.g. [14, 15]). In these models there was no direct gene-to-gene coupling between coevolving species, but the first species consisted of a population of candidate solutions to

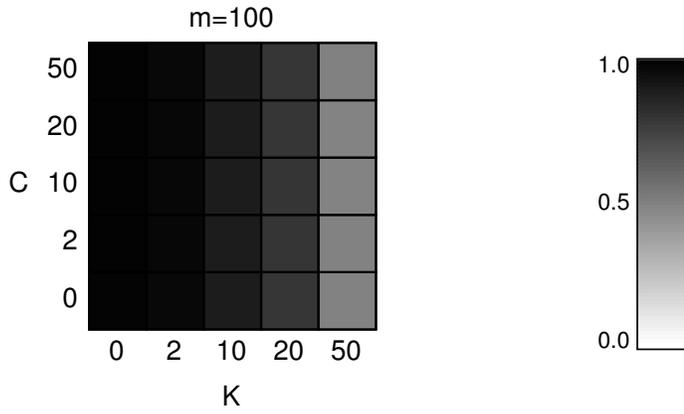


Figure 3: A grey-scale plot of r_1 for the different K and C combinations for $m=100$. The scale on the right shows the grey-scale values.

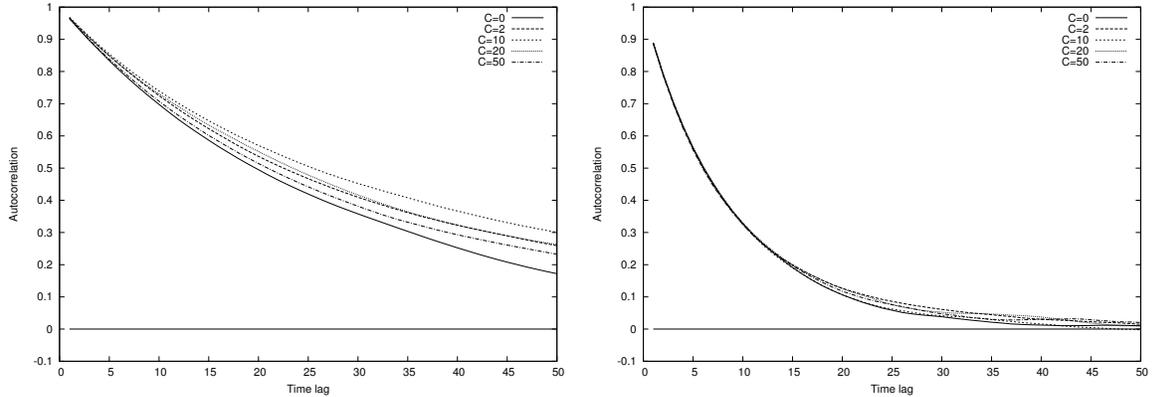


Figure 4: The autocorrelations for $m=100$ for $K=2$ (left) and $K=10$ (right).

the problem, and the second species of a population of test cases. The fitness of the individuals in the first population (the solutions) was then determined by the number of test cases from the second population they are able to solve correctly, and the fitness of the test cases was determined by the number of solutions they are able to “mislead” (i.e., make the solutions perform poorly). In other words, the two species formed each others (changing) environment, and are thus environmentally coupled but not epistatically (sometimes referred to as “host-parasite” coevolution).

Although there is evidence that in some cases it is not only, or not even largely, the coevolutionary aspect that gave rise to an increase in the quality of the solutions found [16], it is clear that this type of coevolution can give an advantage over evolution in a static environment. Perhaps this more “loose” form of coupling (i.e., environmentally instead of epistatically) somehow corresponds to the above situation of low K and slow coevolutionary rate, thus increasing the overall correlation length of the landscape, enabling evolution to find slightly better solutions. It would certainly be worth investigating this question in more detail.

As the plot on the right ($K=10$) shows, for larger values of K this effect has disappeared again though, and indeed, for intermediate and large values of K a slow coevolutionary rate does not have any effect at all on the correlation structure of the fitness landscape.

5.2.2 Intermediate rate of coevolution

Next, consider the case of an intermediate rate of coevolution, or intermediate values of m . Figure 5 shows grey-scale plots of the main autocorrelation values (r_1) for $m=10$ (left) and $m=5$ (right).

These plots show that there is only a minor effect on the main autocorrelation values, and only for smaller values of K . In the first two or three columns, the grey-scale for higher values of C becomes

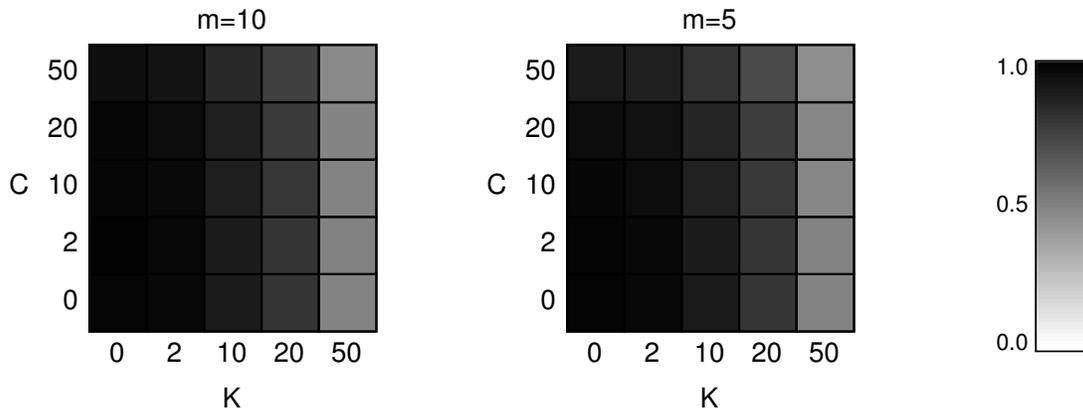


Figure 5: A grey-scale plot of r_1 for the different K and C combinations for $m=10$ (left) and $m=5$ (right). The scale on the right again shows the grey-scale values.

slightly lighter. To have a more detailed view again, figure 6 shows the autocorrelations for $m=10$ for $K=2$ (left) and $K=10$ (right).

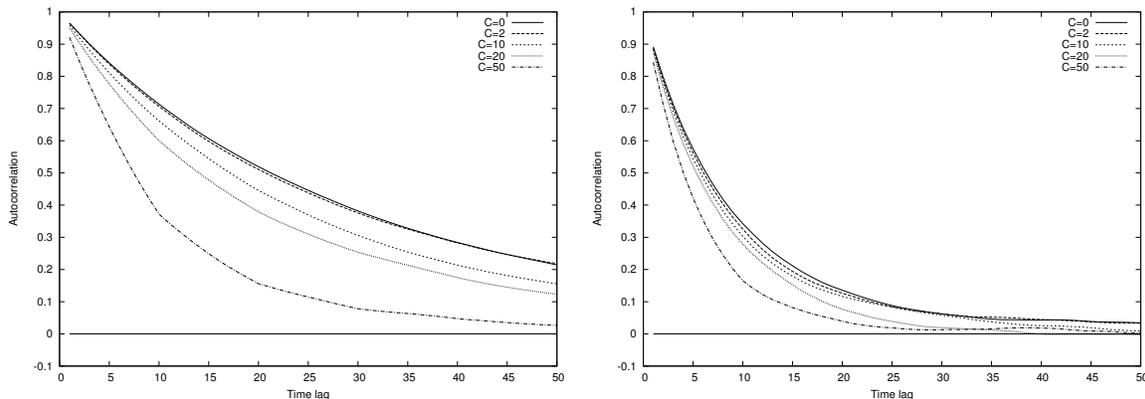


Figure 6: The autocorrelations for $m=10$ for $K=2$ (left) and $K=10$ (right).

From these plots, it is clear that an intermediate rate of coevolution causes the overall correlation length to decrease compared to the static case, but only significantly so for values of $C > K$. In other words, for an intermediate rate of coevolution the internal epistatic interactions are the dominant ones, and only when the level of external epistasis exceeds the internal level does the coupling start to have an influence on the overall correlation structure of the landscape. To illustrate this point even more, figure 7 shows the autocorrelations for $K=20$ for $m=10$ (left) and $m=5$ (right). As the plot on the left ($m=10$) shows, there is hardly any appreciable difference at all for the different values of C , and indeed, for larger values of K the rate of coevolution has to become increasingly faster ($m=5$, right plot) to have any effect on the overall correlation structure, and then again only significantly so for $C > K$.

5.2.3 Fast rate of coevolution

Finally, consider a fast rate of coevolution, in particular when the coevolving species evolve equally fast. Figure 8 shows a grey-scale plot for the main autocorrelations (r_1) for $m=1$. Here, it is clear that the main autocorrelation value decreases quickly with an increasing value of C , and there is a major effect on the correlation of the landscape in all cases.

To show this effect even more clearly, figure 9 shows the autocorrelations for $m=1$ for $K=2$ (left) and $K=20$ (right). The correlation length decreases significantly and ever more drastically with larger values of C , in this case also for $C < K$. In other words, any level of coupling directly influences the

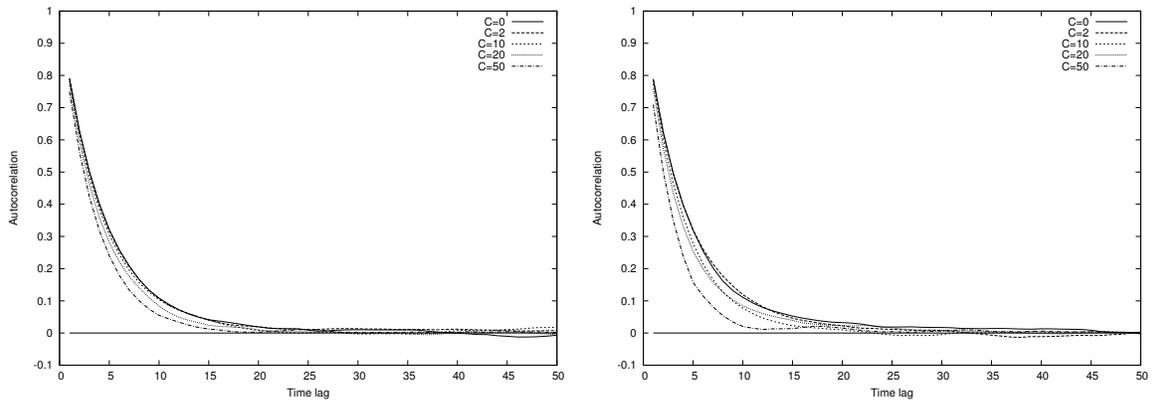


Figure 7: The autocorrelations for $K=20$ for $m=10$ (left) and $m=5$ (right).

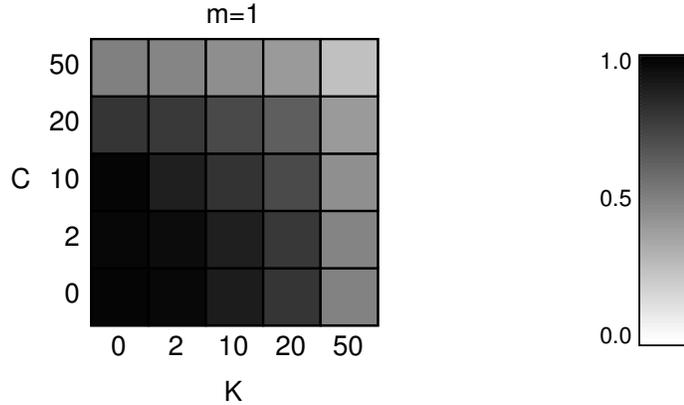


Figure 8: A grey-scale plot of r_1 for the different K and C combinations for $m=1$. The scale on the right again shows the grey-scale values.

correlation structure of the landscape with a fast rate of coevolution, regardless of the amount of internal epistasis.

To conclude this part of the analysis, we note that for larger values of K (e.g. $K=50$; results not shown), there is no appreciable effect, not even for a fast rate of coevolution. This is largely due to the fact that for larger values of K the correlation of the landscape is already so low (see figures 1 and 2), coevolution has very little room left to change this amount of correlation, neither in a positive nor in a negative way. In other words, for a high amount of internal epistasis, the landscapes are already highly rugged and uncorrelated. Adding external couplings is not going to make any significant difference.

5.3 Partial autocorrelations on coupled landscapes

As with regular NK landscapes, in almost all cases considered here only the first partial autocorrelations a_1 were significant, indicating that an AR(1) model is the most appropriate for expressing the correlation structure of these coupled landscapes (see [6] for a detailed explanation of this argument). However, there were a few interesting exceptions.

Figure 10 shows the partial autocorrelations for three different cases: a regular NK landscape with $K=0$, and two NKC landscapes with $m=5$ and $K=20$, $C=10$ and $K=2$, $C=50$, respectively. As the plot shows, in the first two cases indeed only the first partial autocorrelation is significant (although off the scale on this plot), and the others all fall well within the plus or minus two times standard error range (indicated by the horizontal dashed lines). However, for the third case ($K=2$, $C=50$), an interesting pattern can be observed, where every fifth partial autocorrelation seems to be significantly larger than the rest, with the effect slowly tapering off for larger time lags. Clearly, this reflects the

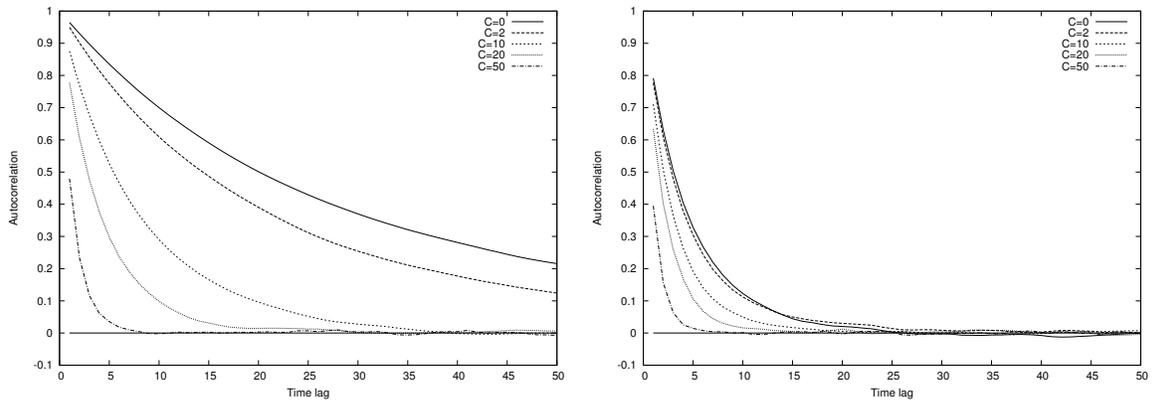


Figure 9: The autocorrelations for $m=1$ for $K=2$ (left) and $K=20$ (right).

fact that $m=5$ was used in this case, i.e., the second (coevolving) species was mutated every five time steps. Several other low to intermediate K and high C cases show a similar pattern in the partial autocorrelations.

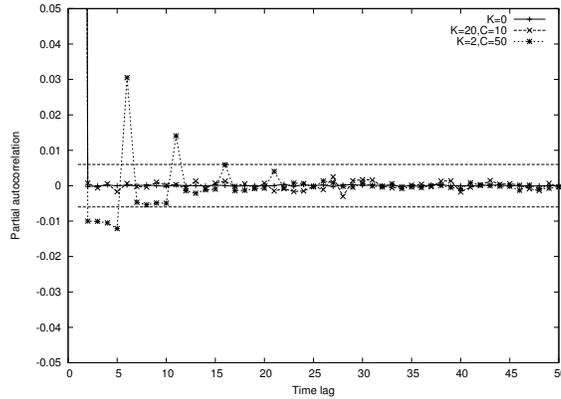


Figure 10: The partial autocorrelations for some landscapes for $m=5$ and various values of K and C (see legend).

This observation seems to indicate that for highly coupled species with low to intermediate internal epistasis, the partial autocorrelations can say something about their relative rates of evolution. Suppose, for example, that the partial autocorrelations of the landscape of the first species show a period three pattern (similar to the period five pattern observed above), then this could indicate that the rate of (accepted) mutations in the second species is three times slower than that of the first species. In other words, there is only one step (accepted mutation) in the second species for every three steps (accepted mutations) in the first species. Perhaps this could indicate the difference between a “younger” (recently evolved) and an “older” species, where the younger species still evolves at a faster rate than the older species in which fewer and fewer mutations will lead to an increase in fitness. This forms an interesting hypothesis that could, for example, be tested in other simple models of coevolution.

As the results in this section clearly show, the correlation structure of coupled fitness landscapes depends very much on the relative rate of evolution of the coevolving species. In other words, there are multiple time scales involved, as is the case in most evolutionary processes. However, in the case of static fitness landscapes, these different time scales are often “ignored”, and it is argued that whenever a species makes an adaptive move, the whole population converges very quickly towards the new location in the fitness landscape, and everything continues as usual again on a larger time scale. This is the equivalent of a slow rate of coevolution in our analysis. But, obviously, the situation becomes different when we consider an intermediate or fast rate of coevolution, and thus our method

and analysis provide at least one way of distinguishing different time scales in evolution, taking them into account explicitly.

6 Conclusions and Future Directions

The main conclusions of our analysis can be summarized as follows:

- First of all, the correlation analysis for fitness landscapes as introduced in [6] can be extended to coupled (i.e., dynamic) fitness landscapes in a simple and meaningful way.
- For a slow rate of coevolution, there is no effect on the correlation structure of the landscape, except for small values of K where the correlation length of the landscape actually increases. This could indicate that a “loose” coupling between species can be mutually beneficial. This was indeed shown in several genetic algorithm studies with environmentally (but not epistatically) coupled species.
- For an intermediate rate of coevolution, there is only an effect when $C > K$, in which case the correlation length of the landscape decreases significantly, with a minor decrease in the main autocorrelation value r_1 . When the level of external epistasis is lower than the level of internal epistasis, it will not have an effect on the overall correlation structure of the landscape.
- For a fast rate of coevolution, any level of external coupling will have a significant negative effect on the correlation length and the main autocorrelation value, except for very large values of K (high level of internal epistasis), in which case the landscape already is highly rugged and uncorrelated anyway.
- In coevolving species with a low level of internal but a high level of external epistasis, the partial autocorrelations can possibly provide insight into the relative rates of evolution of the different species. This provides yet another advantage of the extended correlation analysis as introduced in [6] over the more restricted analyses used elsewhere.
- Finally, the correlation analysis of coupled fitness landscapes as presented here produces meaningful and interesting results that can provide a better understanding of, and lead to testable hypothesis about, coevolution in general.

Several directions for future research and testing of hypotheses were already indicated in previous sections. For example, it would be interesting to investigate whether the landscapes of the environmentally coupled species in some of the genetic algorithm studies mentioned [14, 15] indeed show similarities to the low K , slow rate of coevolution landscapes analyzed here, where the correlation length actually increases compared to the static case. Furthermore, the hypothesis about the partial autocorrelations indicating relative rates of evolution in low K , high C landscapes should be tested in other coupled landscape models. Obviously, it will also be useful to apply to correlation analysis for coupled fitness landscapes presented here to other coevolutionary models, and relate the results to, for example, observations from studies on (co)evolutionary dynamics on those landscapes. And finally, similar to the static case, where much theoretical work is done on the (correlation) structure of fitness landscapes (see e.g. [17, 18, 19, 20, 21]) an actual theory of the correlation structure of coupled fitness landscapes should be developed.

Acknowledgments

WH would like to acknowledge funding from the ACI IMPBIO (French Ministry of Research) while at the LIRMM.

References

- [1] S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones, editor, *Proceedings of the Sixth International Congress on Genetics*, pages 356–366, 1932.

- [2] S. A. Kauffman and S. Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128:11–45, 1987.
- [3] S. A. Kauffman. Adaptation on rugged fitness landscapes. In D. Stein, editor, *Lectures in the Sciences of Complexity*, pages 527–618. Addison-Wesley, 1989.
- [4] S. A. Kauffman. *The Origins of Order*. Oxford University Press, 1993.
- [5] S. A. Kauffman and S. Johnson. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology*, 149:467–505, 1991.
- [6] W. Hordijk. A measure of landscapes. *Evolutionary Computation*, 4(4):335–360, 1996.
- [7] E. D. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336, 1990.
- [8] G. E. P. Box and G. M. Jenkins. *Time Series Analysis, Forecasting and Control*. Holden Day, 1970.
- [9] C. W. J. Granger and P. Newbold. *Forecasting Economic Time Series*. Academic Press, 2nd edition, 1986.
- [10] B. Manderick, M. de Weger, and P. Spiessens. The Genetic Algorithm and the Structure of the Fitness Landscape. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 143–150. Morgan Kaufmann, 1991.
- [11] M. Lipsitch. Adaptation on Rugged Landscapes Generated by Iterated Local Interactions of Neighboring Genes. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 128–135. Morgan Kaufmann, 1991.
- [12] W. Hordijk. Correlation analysis of the synchronizing-ca landscape. *Physica D*, 107:225–264, 1997.
- [13] L. Altenberg. Nk fitness landscapes. In T. Back, D. Fogel, and Z. Michalewicz, editors, *The Handbook of Evolutionary Computation*. Oxford University Press, 1997. Section B2.7.2.
- [14] W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228–234, 1990.
- [15] H. Juillé and J. B. Pollack. Coevolutionary learning and the design of complex systems. *Advances in Complex Systems*, 2(4):371–394, 2000.
- [16] J. Werfel, M. Mitchell, and J. P. Crutchfield. Resource sharing and coevolution in evolving cellular automata. *IEEE Transactions on Evolutionary Computation*, 4(4):388–393, 2000.
- [17] T. C. Jones. *Evolutionary Algorithms, Fitness Landscapes and Search*. PhD thesis, University of New Mexico, Albuquerque NM, USA, March 1995.
- [18] P. F. Stadler. Towards a theory of landscapes. In R. López-Pena, R. Capovilla, R. García-Pelayo, H. Waelbroeck, and F. Zertuche, editors, *Complex Systems and Binary Networks*, pages 77–163. Springer Verlag, 1995.
- [19] P. F. Stadler. Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20:1–45, 1996.
- [20] R. Happel and P. F. Stadler. Canonical approximation of fitness landscapes. *Complexity*, 2:53–58, 1996.
- [21] W. Hordijk and P. F. Stadler. Amplitude spectra of fitness landscapes. *Advances in Complex Systems*, 1(1):39–66, 1998.