

Artur Matos, Reiji Suzuki and Takaya Arita, "Heterochrony and Evolvability in Neural Network Development", *Artificial Life and Robotics*, Vol. 11, No. 2, pp. 175-182, 2007.

(Draft Version)

Abstract

Recent studies in evolutionary computation have focused on using developmental processes together with genetic algorithms in order to achieve more complex designs. Although several models have been proposed, their growth dynamics, and their interactions with evolutionary algorithms are still poorly understood. One particularly neglected concept in artificial developmental systems is *heterochrony* — how evolution affects development by changing the timing and rate of developmental events. In this paper we attempt to address this issue by analyzing heterochronic changes in a well known artificial developmental model — the cellular encoding model — by using an heterochrony framework by Alberch *et al.* We have conducted experiments by evolving networks to solve a boolean problem, and analyzed heterochronic changes in both successful and unsuccessful runs. Our findings show that the cellular encoding model, due to its properties, strongly affects the developmental dynamics and the heterochronic changes that occur during evolution. Our experiments also show that hypermorphic changes (a kind of heterochronic occurrence) lead to greater evolvability in successful runs.

1 Introduction

Recently there has been an increasing interest within the EC (Evolutionary Computation) community in simulating developmental processes alongside evolution.¹ This approach has already proved to be fruitful, allowing evolutionary algorithms (EAs) to generate more complex designs than traditional approaches, and it has been applied to a wide range of domains, including, among others, neural networks² and artificial creatures.³ Although several of these models have been proposed, there is still no throughout understanding on how they work, specially on how evolution and development interact. One particularly important area, and that we will address in this paper, is to try to understand how the EA shapes the individuals by rearranging the underlying developmental events.

In evolutionary and developmental biology, this change in the rate, timing and order of developmental events caused by evolution is generally known as *heterochrony*.⁴ Heterochrony is a well observed phenomena and prevalent in the evolution of species. A well known example can be found in the Mexican axolotl salamander: most salamander species have two distinct stages of development, a larval and an adult stage. However, the Mexican axolotl does not undergo metamorphosis, and achieves sexual maturity in what would still be considered a larval form. Therefore, from an evolutionary point of view, we can say that evolution shaped the axolotl species by “slowing down” their ancestor’s development. Because examples like this are so common in nature, it is speculated that heterochrony is one of the major factors in the evolution of more complex taxa.

Unfortunately, and despite of this, EC studies on heterochrony are still lacking. A recent review paper on artificial developmental models¹ identified heterochrony as one of the important dimensions to pursue in research, but there are still very few published results on the topic:⁵ described heterochronic occurrences in a developmental model based on genetic regulatory networks (GRNs). The author evolved neural networks for solving a food foraging problem, and then compared developmental events between ancestor and descendant networks. By proceeding this way, he was able to identify several different kinds of heterochrony, for instance, occurring in cell division events and during axon growth.³ also used a GRN-based model, but for evolving artificial creatures, with coevolved morphologies and neural networks. The author evolved individuals for a locomotion task, and then performed mutation experiments to analyze the evolved GRNs. He then observed that mutation in some individuals caused morphological units to appear earlier or later compared to the original individual.

Although these previous studies show that heterochrony does indeed occur in artificial developmental models, there are still several questions remaining unanswered. First, it is still not clear to what extent these models support heterochrony, or if they direct evolution to certain kinds of heterochronic occurrences more than others. A second, and perhaps even more important question is to understand how heterochrony relates to EA performance: for instance, can we expect certain kinds of heterochronic occurrences to be more conducive to evolvability than others? As a starting point for the first question, in a previous paper,⁶ we have applied a heterochrony framework to a well known developmental model — the cellular encoding model developed by Gruau.² This framework, by

Alberch *et al.*,⁷ offers both a precise terminology and methodology to study heterochronic phenomena in living systems. In this paper, we extend our previous results, and also attempt to undertake the second question. For this, we have used the cellular encoding model to evolve neural networks to solve a boolean problem, and applied the framework to a large ensemble of networks. This allows us to characterize heterochrony in artificial systems in a more complete way than the previous studies. For the second question, we have applied the framework to both successful and unsuccessful experiments, and compared the heterochronic occurrences between them, to check for any meaningful differences.

2 Alberch *et al.*'s framework

The framework by Alberch *et al.* is widely used in biological systems for analyzing heterochronic phenomena. This framework is based on the measurement and comparison of quantitative traits, for instance, body length, width or height. The traits are measured as development unfolds, yielding growth curves. These growth curves can then be compared between related species for understanding the heterochronic change involved.

The basis for comparison lies on three metrics that can be extracted from the growth curves: α — the time when growth starts, β — the time when growth ends, and K — the growth rate. Comparing these values between species yields the outcomes summarized in Figure 1. For instance, considering only changes in the K parameter, two outcomes are possible: if the descendant would grow faster than the ancestor (K would be larger), the corresponding outcome is *acceleration*. The reverse process — the descendant growing slower than the ancestor — is labeled *neoteny*. Furthermore, changes can be combined on the three parameters, although we didn't consider them in this analysis.

3 Cellular encoding

Cellular encoding is a developmental model originally proposed by Gruau² for evolving neural networks. We have chosen this model because it has proven track for evolving neural networks, for a wide range of problems, including threshold neural networks for boolean problems² and controllers for robot locomotion.⁸ Instead of using GRNs as in⁵ or,³ the cellular encoding model specifies development as a set of graph-rewriting instructions that are evolved directly. The original model, that we follow in this paper, only evolved simple threshold networks: the neurons were threshold neurons with thresholds of either 0 or 1, with the connections between the neurons being either -1 or 1.

Cellular encoding defines a set of commands for operating on graphs, that changes the graph as development unfolds. In its original description, development is described as a sequence of these commands, grouped in a Genetic Programming (GP) tree. These trees are then used just as standard genotypes in a GP system. Each network starts as a single neuron, with a pointer pointing to the root of the tree. Development on each neuron proceeds

by executing sequentially the nodes with the developmental commands. There are different kinds of commands, including commands for dividing cells, creating new connections, setting thresholds and so on. For instance, the **PAR** instruction, when applied to an existing neuron, creates a new neuron and copies the connections (both the input and output connections) from the original neuron to the new one. **PAR** nodes contain two children, that are inherited by each neuron (the original and the new one), allowing for cell differentiation to occur. Development in the network occurs in a parallel fashion: on each time step each neuron executes the command pointed by its register in the tree and moves to the following leaf. The developmental process is over when all the neurons have reached their final leaf node in the tree. A summary of the commands is shown in table 1.

4 Experiments

In his original experiments, Gruau evolved networks only for solving boolean problems, for instance, the odd-parity problem and the symmetry problem. In this paper, we decided to use a different boolean problem: the function we are trying to optimize has 3 inputs and 1 output, with the inputs ranging from 000 to 111. All the inputs in this function produce the output 0 except for two entries corresponding to inputs 011 and 100, where the output is 1. Our experiments have shown that this problem is particularly difficult to solve, and therefore could be used as a good problem for checking relationships between heterochrony and evolvability.

The fitness function used consisted of two parts: the first accounting for the number of right output values, and the other part rewarding networks with the right number of input and output neurons. The first part just computes the number of correct output values, ranging from 0 to 8. The second part returns the number of input neurons (or 3, whichever is smaller), and the number of output neurons (or 1, whichever is smaller). This latter part was deemed necessary to ensure that the problem could be solved at all. They are then weighted (with 0.75 for the number of right output values, and 0.25 for the number of right input/output neurons), and normalized between the range $[0, 1]$ with 1 as the best fitness. Furthermore, because tree depth tended to increase rapidly during evolution (tree bloat), we have imposed an upper limit of 30 neurons in the networks; all the networks that exceed this limit were assigned a fitness value of 0.

We used a GP system, with both population size and the number of generations set to 300. Tournament selection was used, with a tournament size of 7. Crossover was not used, and all the individuals were mutated with a 70% probability. The mutation operator follows Koza's GP original description: if a mutation occurs, first a node is selected at random in the GP tree; The node is then replaced with a new random subtree, with a maximum depth of 5.

We have conducted several evolutionary runs, all with different random seeds. Because of the problem difficulty, several of the runs never found any optimal network, and converged into a local optima. For analysis, we kept 80 runs of all the ones conducted: 40 where the optimum was reached, and 40 that converged to a local optima.

5 Analysis

5.1 (α, β, K) dynamics

All of the successful runs exhibited similar behavior. On average, it took 110.7 generations to reach the optimum. A fitness graph of a typical run can be seen in Figure 2, and an optimal network in Figure 3.

In order to apply the Alberch *et al* framework to the model we needed to choose suitable traits for analysis. We decided to consider the number of nodes and the average network degree (considering both incoming and outgoing connections), as they are the essential measures related to the network topology. We also decided to analyze the fitness dynamics, and treated the intermediate fitness values — how the fitness changes during development — as another trait in the framework. By proceeding this way, we could then check if the fitness dynamics played an important role in evolution, although for evaluation purposes only the fitness value of the last developmental step is considered. To get an idea how these traits change with development, Figure 4 shows growth curves of these 3 traits, in two different individuals. The dynamics on the right are from the best network in Figure 3.

Furthermore, to apply the framework, we need to extract the three parameters — α , β and K — from the growth curves. The framework itself does not define a method for doing this, and several methods have been proposed in the literature. In biological studies, one common method is to fit the data to a growth model (for instance, the Von Bertalanffy growth curve) by using non-linear regression, and then extract the parameters from the fitted equation. This works well for biological data because they tend follow well-known patterns, and there are several sensible growth models available in the literature. In contrast, as shown in Figure 4, our growth curves tend to be rather irregular, so this approach is not feasible. Therefore we decided to use a simpler approach: for our growth curves, we defined α and β as the developmental time where growth effectively starts and stops in the data, that is, where changes in the values occur for the first and last time during development. As for K , we defined it as the average growth rate. Figure 5 shows the evolution of these three parameters in one typical, successful run. This example refers to the best lineage, that is, all of the individuals, starting from the first generation, that gave rise to the best individual in the last generation.

In this example, we can see that beta increased and K decreased through generations. This heterochronic change corresponds to the combination of hypermorphosis and neoteny in Figure 1. One interesting point is that α didn't change at all during evolution, for any of the 3 traits, and in any of the conducted runs. The reasons beyond this are two: the first reason is that the cellular encoding model represents the developmental events sequentially in the genotype, following tree order. Although the mutation operator does choose a node randomly in the tree, successful mutations tend to target nodes at a greater depth, because these tend to create smaller changes in the phenotype, and allow evolution to occur gradually. A change in the top-level node of the tree basically amounts to replace the original individual with a completely unrelated one, so it does not take into account any of the previously found solutions. The second reason is that in the cellular encoding model, almost all of the instructions must always

change the network in some way; therefore it is difficult for the model to generate delays in the timing of events.

5.2 Correlation analysis

We decided further to check if there was any significant correlation in the dynamics, by applying the Pearson's correlation coefficient to the data:

$$\text{cor}(X, Y) = \frac{\sum (X_i - \bar{X})(Y_i - \bar{Y})}{(n - 1)S_X S_Y} \quad (1)$$

with X_i and Y_i the parameter values of X and Y at generation i , \bar{X} and \bar{Y} as the mean values of X and Y , and S_X and S_Y as their standard deviation. A coefficient of 1 indicates a perfect correlation, and -1 a negative perfect correlation.

We have applied this coefficient between individual parameters in the same trait (for instance between β and K of the number of nodes trait), as well as for the same parameters in different traits (for instance, between β in the number of nodes trait and β in the average degree trait). Because α never changes during the runs, α was not considered in this analysis. We have also computed the coefficients separately for successful and unsuccessful runs. The results are shown in Table 2. This data refers to the best lineages only, as in the example shown in Figure 5.

The correlation table shows how the cellular encoding model further constrains the developmental dynamics. To understand it better, first it is necessary to explain how the model creates connections in the networks: in the cellular encoding model, there is no explicit command for adding new connections, although there is a command for removing them. The only way to create new connections in the networks is by cell division, that is, by creating new neurons. Therefore the number of nodes trait and average degree cannot evolve independently, and become highly correlated. This can be seen in the correlation coefficient of the number of nodes β and average degree β : in most cases, the growth of the average degree trait stops growing effectively when the number of nodes also stops growing, except when remove connections commands are found after the last cell division command. This seems to rarely happen, however, as it is shown in the table. A similar reason is behind the high correlation between the average degree's β and K : as indicated before, our fitness function imposes an upper limit on the number of neurons that a network may have; this combined with what was explained now effectively defines an upper limit on the average degree as well. This makes β and K negatively correlated, in order for the networks not to exceed this upper limit. This also explains why on the unsuccessful runs this correlation is lower.

The lower rows in the table, concerning correlations between the topological traits and fitness are also important, because they show how the dynamics contribute to the fitness as whole. The first thing that can be observed is that the average degree parameters are more strongly correlated with fitness than the number of nodes. This is to be expected, because the connectivity of the network (the way that the neurons are connected to each other) is more important for solving this problem than the number of nodes. Other important point is that higher β correlations

between the topological traits and fitness seem important to assure successful evolution. Unfortunately, because we have only conducted analysis using this problem so far, at this time it still difficult to say why this occurs.

5.3 Occurrence analysis

Based on the previously computed parameters, we further classified the changes according to the framework, with the results depicted on Table 3. Results for all the three traits were similar, so only the number of nodes trait is displayed. As it is shown in the table, all possible outcomes (with the exception of predisplacement and postdisplacement, as it was explained before) occur with significant values; this, therefore, shows that the cellular encoding model is able to generate most kinds of heterochronic events. We can see that there is a difference in the occurrences among the heterochronic changes. For instance, isomorphosis was the most frequently observed occurrence in successful and unsuccessful runs, neoteny the second most frequent one, but progenesis rarely occurred in both runs. The relative distribution of these occurrences is basically the same between successful and unsuccessful runs, but we can also see that each heterochronic occurrence was slightly less observed in successful runs compared with unsuccessful runs. All these differences are statistically significant.

On average, 54% of the occurrences in Table 3 were real changes (non-isomorphic); this is much lower than what should be expected, considering the mutation probability that was used (70%). This means that selection favors isomorphic changes in the system (that is, neutral mutations), and overall it should be favoring isomorphic changes in the successful runs. One possible reason for this is that there could be an increased pressure for neutral changes after the optimum has been reached. We decided to check if this occurred, by performing the following analysis: on all experiments, we have divided the individuals in the best lineages into two groups: the first group with all the individuals before the optimum has been reached, and the second group with the remaining ones. For the unsuccessful experiments, we considered the highest reached value as the optimum. On almost all the experiments, the number of isomorphic changes increases after the optimum was reached, regardless of being a successful run or not. We have applied a Fisher exact test to each experiment to check for statistical significance, and the difference in ratios was significant in 16 of the 40 successful experiments, and exactly the same number of unsuccessful experiments. We believe that this difference in ratios was significant only for a small number of experiments due to the the small sample size in some of the runs: in some of the experiments the optimum is found rather early in evolution, and this makes the sample size before the optimum very small compared with after the optimum is reached; this causes the statistical test to fail in these experiments, although the number of isomorphic changes almost always increases after the optimum is reached (with the exception of 2 experiments out of the 80). Because there is no significant difference between successful and unsuccessful runs in these cases, we cannot conclude that neutral mutations are concentrated after the optimum is reached in the successful runs. Therefore, although successful runs have an increased rate of isomorphic occurrences, they are evenly distributed during evolution.

Extending this statistical analysis to cover all kinds of occurrences and all the experiments is difficult, and it will

not be attempted in this paper; nevertheless, in table 4, we summarize these results. Table 4 shows the average ratios between each kind of heterochronic occurrence dividing by the total number of occurrences, before and after the optimum was reached. We can see on the table that isomorphosis does indeed occur with more frequency after the optimum is reached, and that the other heterochronic occurrences are less frequently observed after the optimum is reached. This, therefore, supports our previous point.

We have also decided to compute the total changes in fitness that each kind of occurrence has lead to: for instance, for the neoteny case, we have computed all the differences in fitness values due to neotenic occurrences. The results are shown in Table 5. There were no statistical significant differences to be found between successful and unsuccessful runs, but still we can observe differences among the heterochronic occurrences: first, the total increase in fitness is quite different for each heterochronic change, and it is not proportional to their number of occurrences. For instance, hypermorphosis brought the highest increase in fitness in both successful and unsuccessful runs, although it is only the third most frequent occurrence. This is particularly noticeable in the successful runs. These differences between heterochronic types are all highly statistically significant (with p closer to 0). These results partially show that hypermorphosis tends to yield rapid increases in fitness compared to the other occurrences, and therefore promote evolvability. One possibility is that hypermorphosis is increasing the search space used by evolution, and therefore making mutations more efficient. Because hypermorphic changes correspond to increases in developmental time (β), this also increases the range, in developmental time, where mutations may occur. This, in turn, increases the range of phenotypes reachable by mutation, allowing the search space to be more efficiently explored.

6 Conclusion

In this paper we have shown that the framework by Alberch *et al* is a valid method for studying heterochrony in artificial systems, by applying it to a typical artificial developmental model, in this case the cellular encoding model. Our results show that this model constrains developmental dynamics at least in two ways: first, because the GP trees are executed sequentially during development, predisplacement and postdisplacement events are very unlikely to occur; second, because the cellular encoding model can only create new connections by adding new neurons, development of the topological traits are highly correlated.

Our results concerning heterochrony and evolvability are summarized as follows: first, our results point out that stronger correlations between the way fitness changes during development and the topological dynamics may be important for evolvability. Second, isomorphosis is the most observed change in both successful and unsuccessful runs, and much more frequent than any other kind of heterochrony. Third, successful runs had more isomorphic occurrences than unsuccessful runs, and they are evenly distributed during evolution. Fourth, hypermorphosis was the heterochronic occurrence contributing more to increase fitness, and this should be due to it increasing the search

space used by evolution. However, it is still necessary to apply this framework to other tasks to see if these results generalize to other problems. Future work will focus on these topics, and also extend the current analysis to different kinds of models.

References

1. Stanley, K. O.; Mikkulainen, R. *Artificial Life* **2003**, *9*, 93-130.
2. Gruau, F. *Neural Network Synthesis using Cellular Encoding and the Genetic Algorithm*, Thesis, Laboratoire de la Informatique du Parallelisme, Ecole Normale Supérieure de Lyon, 1994.
3. Bongard, J. Evolving modular genetic regulatory networks. In *Proceedings of The IEEE 2002 Congress on Evolutionary Computation*; IEEE Press: Piscataway, NJ, USA, 2002.
4. Klingenberg, C. P. *Biological Reviews* **1997**, *73*, 79-123.
5. Cangelosi, A. Heterochrony and Adaptation in Developing Neural Networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Vol. 2; Morgan Kaufmann: 1999.
6. Matos, A.; Suzuki, R.; Arita, T. Evolution of Development and Heterochrony in Artificial Neural Networks. In *Proceedings of the Tenth International Symposium on Artificial Life and Robotics*; Beppu, Oita, Japan, 2005.
7. Alberch, P.; Gould, S. J.; Oster, G. F.; Wake, D. B. *Paleobiology* **1979**, *5*, 296-317.
8. Gruau, F. *Adaptive Behaviour* **1995**, *3*, 151-183.

List of Figures

1	The formalism of Alberch <i>et al.</i> A trait measure is plotted against developmental time in the X axis. The solid line plotted from α to β represents the growth curve for the ancestor, while the remaining ones possible heterochronic outcomes for the descendant.	13
2	Fitness graph of a typical run.	14
3	An example of an optimal network, found during our evolutionary runs.	15
4	Growth curves of the analyzed traits. These correspond to the best lineage taken from Figure 2. . .	16
5	Evolution of the α , β and K parameters in the best lineage, in one of the conducted runs. Top row: number of nodes; middle row: average degree; bottom row: fitness.	17

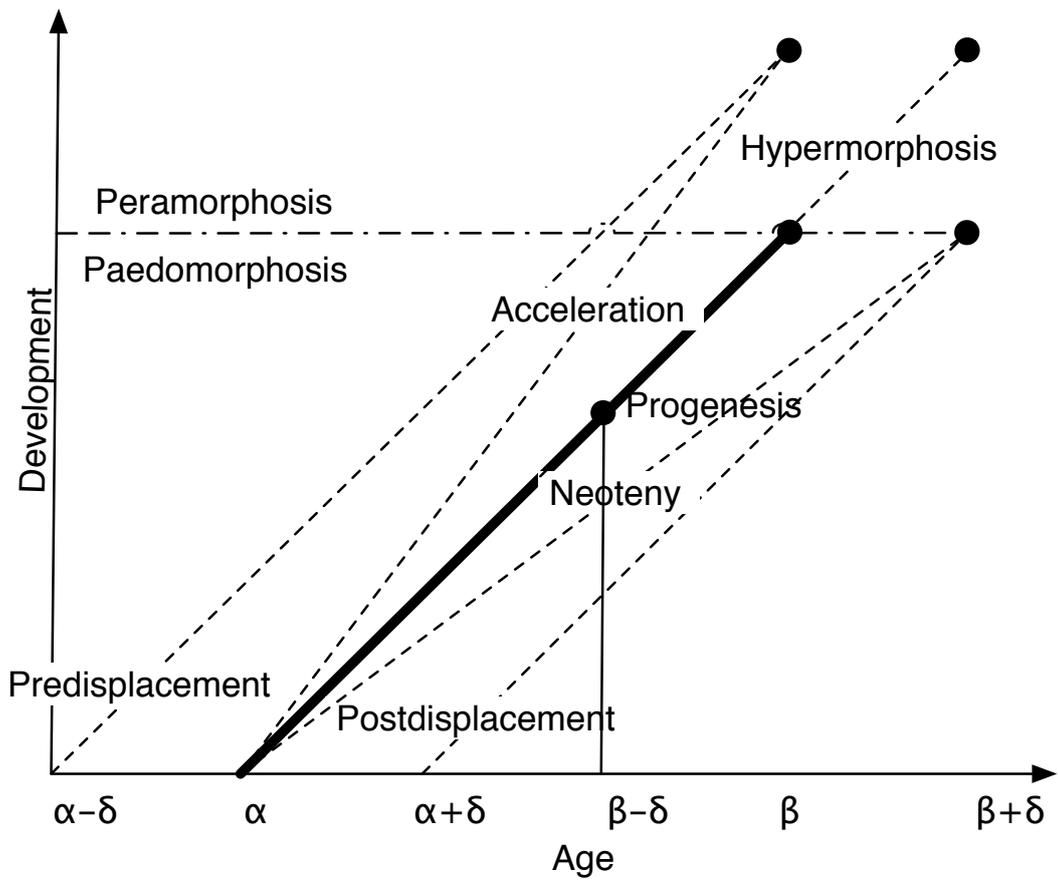


Figure 1: The formalism of Alberch *et al.* A trait measure is plotted against developmental time in the X axis. The solid line plotted from α to β represents the growth curve for the ancestor, while the remaining ones possible heterochronic outcomes for the descendant.

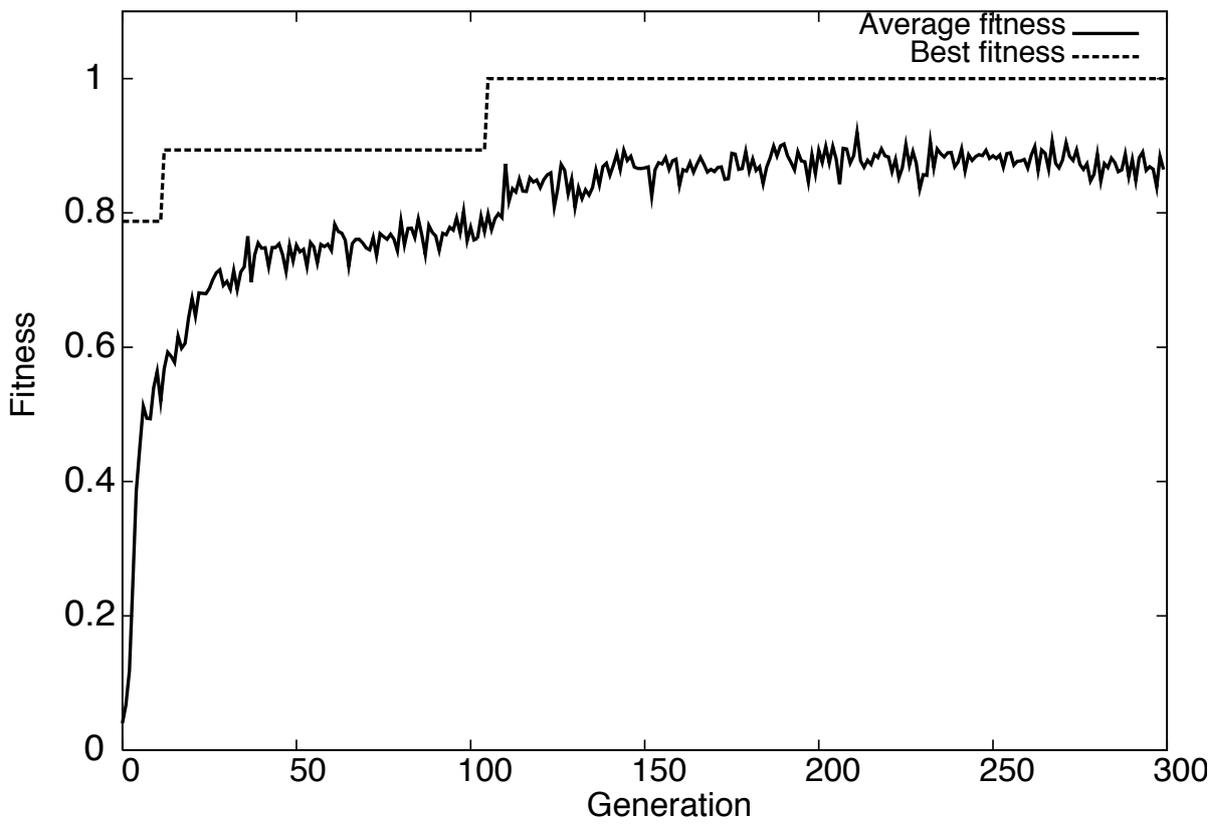


Figure 2: Fitness graph of a typical run.

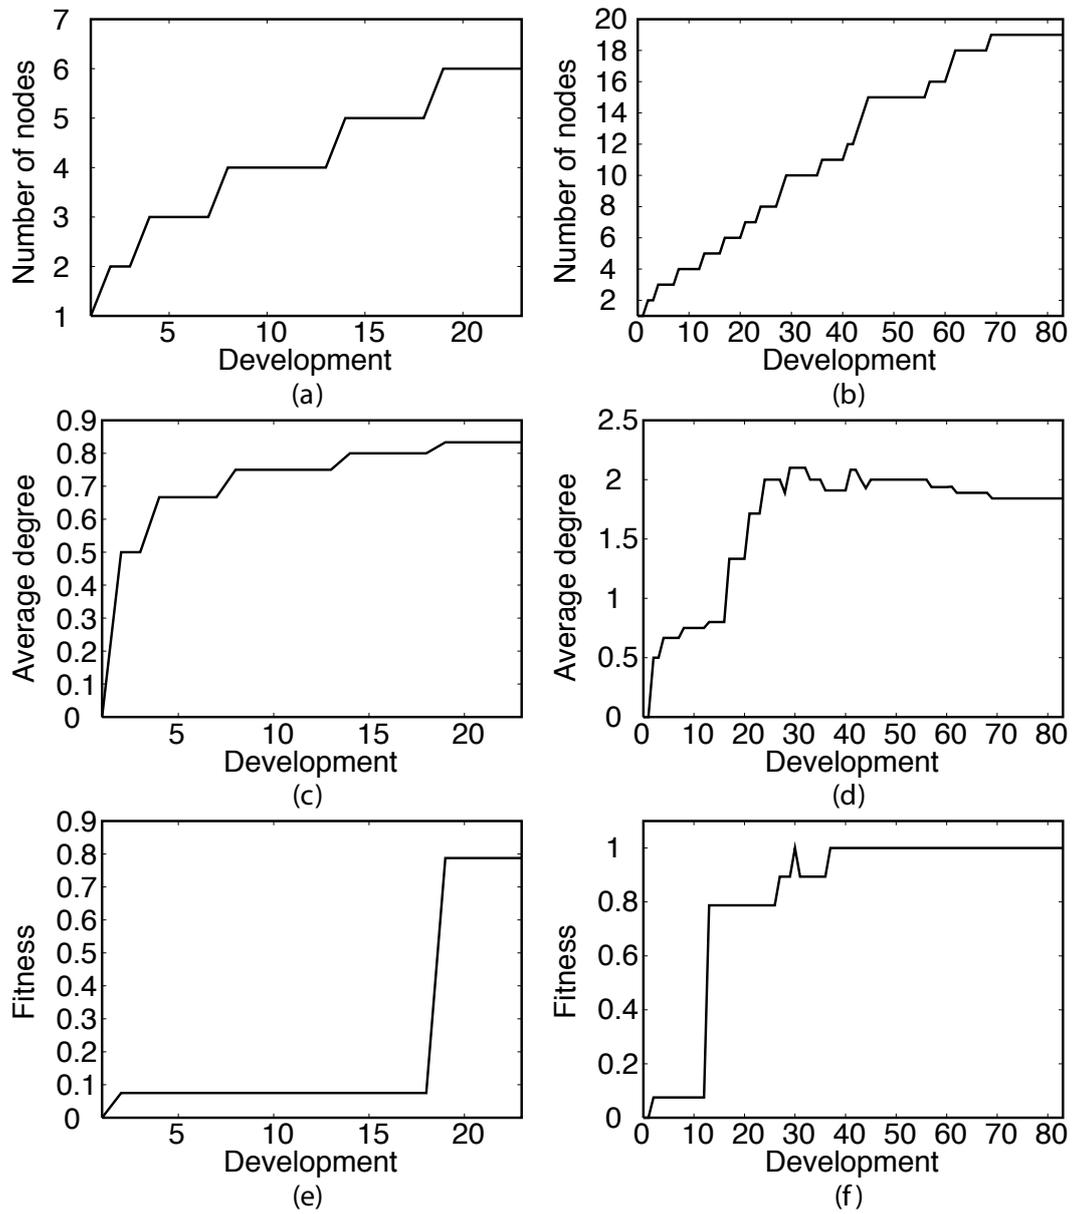


Figure 4: Growth curves of the analyzed traits. These correspond to the best lineage taken from Figure 2.

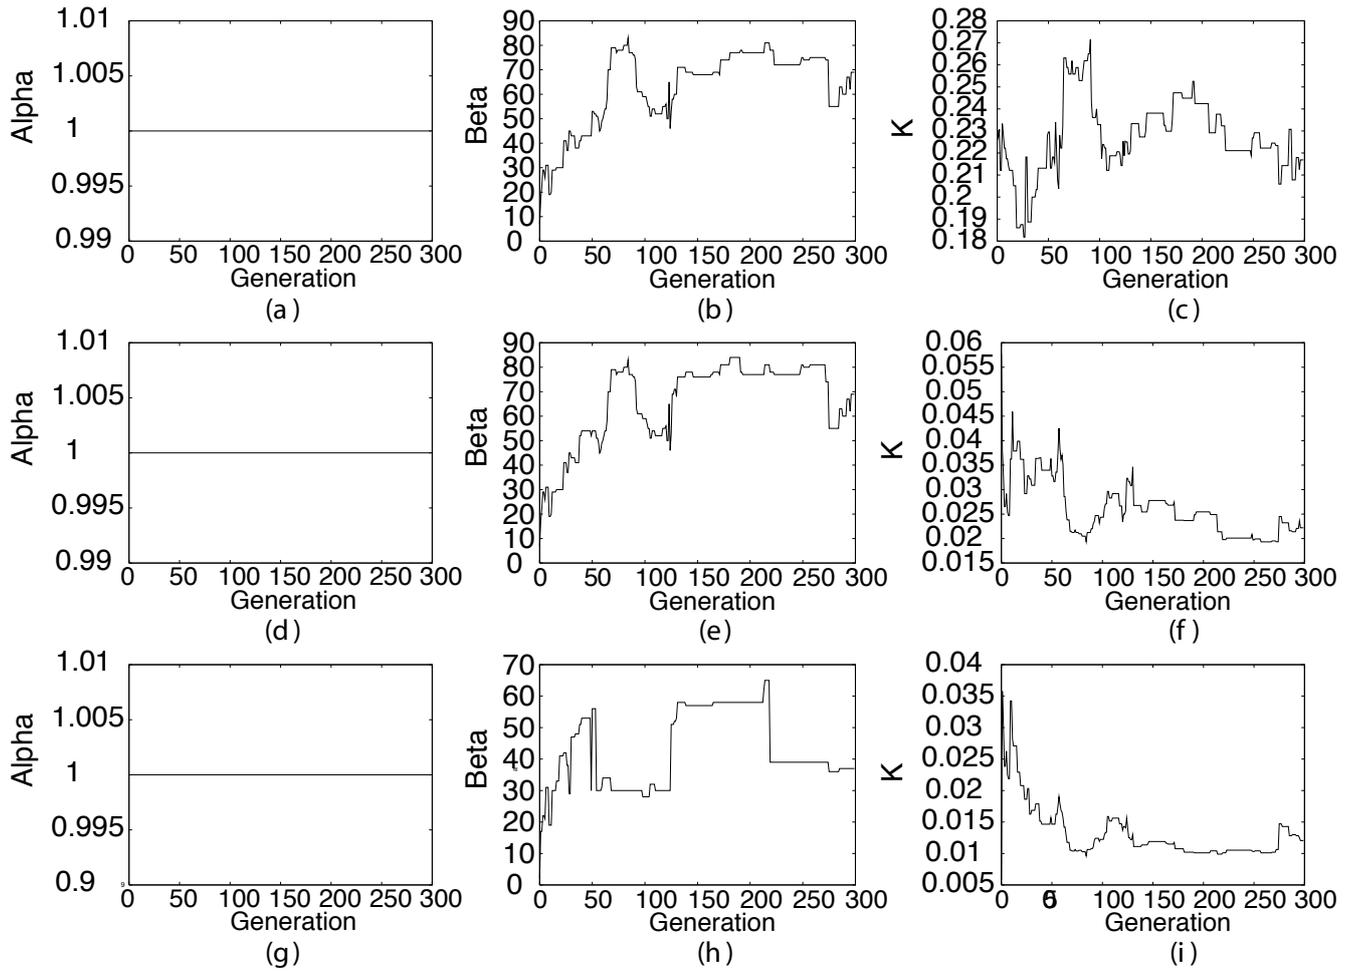


Figure 5: Evolution of the α , β and K parameters in the best lineage, in one of the conducted runs. Top row: number of nodes; middle row: average degree; bottom row: fitness.

List of Tables

1	Cellular encoding commands	19
2	Correlation table for the heterochronic parameters. Entries in gray represent values where the differences between successful and unsuccessful runs showed a strong statistical difference (at the $p = 0.05$ level). p represents the p-value computed by applying t-tests on each set of data. NN - number of nodes; AD - average degree; F - fitness.	20
3	Classification of heterochrony occurrences in all runs for the number of nodes trait. p represents the p-value computed by applying t-tests on each set of data. I - isomorphosis; A - acceleration; N - neoteny; P - progenesis; HM - hypermorphosis.	21
4	Ratios of heterochronic occurrences before and after the optimum was reached, for the number of nodes trait.	22
5	Total increase in fitness due to each heterochronic type. p represents the p-value computed by applying t-tests on each set of data.	23

Table 1: Cellular encoding commands

Command	Description
PAR	Parallel division: Divides a neuron into two, with the incoming and outgoing connections replicated for the new neuron.
SEQ	Sequential division: Divides a neuron into two, the first neuron inherits the incoming connections, while the second neurons inherits the outgoing connections. A single connection is added to the two.
DECBIAS	Sets the neuron's threshold to 1.
INCBIAS	Sets the neuron's threshold to 0.
DECLR	Decreases the link register by 1. The link register is an internal variable in each neuron that points to the current incoming connection that is being manipulated.
INCLR	Increases the link register by 1.
VAL+	Sets the weight of the current incoming connection to 1.
VAL-	Sets the weight of the current incoming connection to -1.
CUT	Removes the current incoming connection.

Table 2: Correlation table for the heterochronic parameters. Entries in gray represent values where the differences between successful and unsuccessful runs showed a strong statistical difference (at the $p = 0.05$ level). p represents the p-value computed by applying t-tests on each set of data. NN - number of nodes; AD - average degree; F - fitness.

Correlation	Successful		Unsuccessful		p
	Mean	SD	Mean	SD	
Same Trait:					
<i>($\beta - K$)</i>					
NN	$-4.153 \cdot 10^{-03}$	$4.084 \cdot 10^{-01}$	$1.912 \cdot 10^{-01}$	$3.267 \cdot 10^{-01}$	$2.063 \cdot 10^{-02}$
AD	$-7.853 \cdot 10^{-01}$	$1.125 \cdot 10^{-01}$	$-7.071 \cdot 10^{-01}$	$1.929 \cdot 10^{-01}$	$2.975 \cdot 10^{-02}$
F	$-5.261 \cdot 10^{-01}$	$2.518 \cdot 10^{-01}$	$-2.845 \cdot 10^{-01}$	$3.133 \cdot 10^{-01}$	$2.840 \cdot 10^{-04}$
Trait-Trait:					
NN - AD					
$\beta - \beta$	$9.793 \cdot 10^{-01}$	$1.544 \cdot 10^{-02}$	$9.672 \cdot 10^{-01}$	$2.913 \cdot 10^{-02}$	$2.297 \cdot 10^{-02}$
K - K	$1.425 \cdot 10^{-01}$	$3.333 \cdot 10^{-01}$	$6.741 \cdot 10^{-02}$	$3.069 \cdot 10^{-01}$	$2.977 \cdot 10^{-01}$
NN - F					
$\beta - \beta$	$5.868 \cdot 10^{-01}$	$2.921 \cdot 10^{-01}$	$2.617 \cdot 10^{-01}$	$3.589 \cdot 10^{-01}$	$2.903 \cdot 10^{-05}$
K - K	$9.159 \cdot 10^{-02}$	$3.598 \cdot 10^{-01}$	$-7.530 \cdot 10^{-02}$	$3.011 \cdot 10^{-01}$	$2.729 \cdot 10^{-02}$
AD - F					
$\beta - \beta$	$6.040 \cdot 10^{-01}$	$2.912 \cdot 10^{-01}$	$2.689 \cdot 10^{-01}$	$3.678 \cdot 10^{-01}$	$2.197 \cdot 10^{-05}$
K - K	$7.640 \cdot 10^{-01}$	$1.566 \cdot 10^{-01}$	$7.559 \cdot 10^{-01}$	$2.173 \cdot 10^{-01}$	$8.494 \cdot 10^{-01}$

Table 3: Classification of heterochrony occurrences in all runs for the number of nodes trait. p represents the p-value computed by applying t-tests on each set of data. I - isomorphosis; A - acceleration; N - neoteny; P - progenesis; HM - hypermorphosis.

Trait	Successful		Unsuccessful		p
	Mean	SD	Mean	SD	
I	$1.822 \cdot 10^{+02}$	$2.544 \cdot 10^{+01}$	$1.695 \cdot 10^{+02}$	$1.970 \cdot 10^{+01}$	$1.485 \cdot 10^{-02}$
A	$5.537 \cdot 10^{+01}$	$1.203 \cdot 10^{+01}$	$6.272 \cdot 10^{+01}$	$9.578 \cdot 10^{+00}$	$3.388 \cdot 10^{-03}$
N	$6.105 \cdot 10^{+01}$	$1.410 \cdot 10^{+01}$	$6.627 \cdot 10^{+01}$	$1.190 \cdot 10^{+01}$	$7.726 \cdot 10^{-02}$
HM	$5.712 \cdot 10^{+01}$	$1.270 \cdot 10^{+01}$	$6.055 \cdot 10^{+01}$	$1.101 \cdot 10^{+01}$	$2.013 \cdot 10^{-01}$
P	$3.990 \cdot 10^{+01}$	$9.139 \cdot 10^{+00}$	$4.492 \cdot 10^{+01}$	$7.494 \cdot 10^{+00}$	$8.764 \cdot 10^{-03}$

Table 4: Ratios of heterochronic occurrences before and after the optimum was reached, for the number of nodes trait.

Occurrence	Successful		Unsuccessful	
	Mean	SD	Mean	SD
I (before)	$3.696 \cdot 10^{-1}$	$9.390 \cdot 10^{-2}$	$2.882 \cdot 10^{-1}$	$1.121 \cdot 10^{-1}$
I (after)	$5.166 \cdot 10^{-1}$	$1.025 \cdot 10^{-1}$	$4.407 \cdot 10^{-1}$	$7.050 \cdot 10^{-2}$
A (before)	$1.506 \cdot 10^{-1}$	$2.295 \cdot 10^{-2}$	$1.771 \cdot 10^{-1}$	$6.580 \cdot 10^{-2}$
A (after)	$1.285 \cdot 10^{-1}$	$3.270 \cdot 10^{-2}$	$1.500 \cdot 10^{-1}$	$2.090 \cdot 10^{-2}$
N (before)	$1.991 \cdot 10^{-1}$	$4.390 \cdot 10^{-2}$	$2.254 \cdot 10^{-1}$	$6.670 \cdot 10^{-2}$
N (after)	$1.355 \cdot 10^{-1}$	$3.680 \cdot 10^{-2}$	$1.573 \cdot 10^{-2}$	$2.320 \cdot 10^{-2}$
HM (before)	$1.854 \cdot 10^{-1}$	$4.710 \cdot 10^{-2}$	$1.945 \cdot 10^{-1}$	$7.030 \cdot 10^{-2}$
HM (after)	$1.266 \cdot 10^{-1}$	$2.780 \cdot 10^{-2}$	$1.422 \cdot 10^{-1}$	$2.290 \cdot 10^{-2}$
P (before)	$9.520 \cdot 10^{-2}$	$3.820 \cdot 10^{-2}$	$1.098 \cdot 10^{-1}$	$5.250 \cdot 10^{-2}$
P (after)	$9.290 \cdot 10^{-2}$	$2.420 \cdot 10^{-2}$	$1.098 \cdot 10^{-1}$	$1.570 \cdot 10^{-2}$

Table 5: Total increase in fitness due to each heterochronic type. p represents the p-value computed by applying t-tests on each set of data.

Occur.	Successful		Unsuccessful		p
	Mean	SD	Mean	SD	
I	$-2.656 \cdot 10^{-03}$	$3.794 \cdot 10^{-02}$	$5.312 \cdot 10^{-03}$	$2.345 \cdot 10^{-02}$	$2.620 \cdot 10^{-01}$
A	$3.101 \cdot 10^{-01}$	$2.807 \cdot 10^{-01}$	$3.525 \cdot 10^{-01}$	$3.449 \cdot 10^{-01}$	$5.488 \cdot 10^{-01}$
N	$2.946 \cdot 10^{-01}$	$2.748 \cdot 10^{-01}$	$2.093 \cdot 10^{-01}$	$2.328 \cdot 10^{-01}$	$1.382 \cdot 10^{-01}$
HM	$4.239 \cdot 10^{-01}$	$3.440 \cdot 10^{-01}$	$3.585 \cdot 10^{-01}$	$3.103 \cdot 10^{-01}$	$3.753 \cdot 10^{-01}$
P	$1.054 \cdot 10^{-01}$	$1.482 \cdot 10^{-01}$	$8.687 \cdot 10^{-02}$	$1.298 \cdot 10^{-01}$	$5.524 \cdot 10^{-01}$