

Evolving in advance: interleaving simulated and physical environments in robot evolution

Twan Goosen (t.goosen@student.ru.nl)
Joris Janssen (jorisjanssen@student.ru.nl)
Pim Haselager (w.haselager@nici.kun.nl)

Abstract

Control structures for physical robots can be evolved in simulated and physical environments. In this study, the interleaving of simulated and physical environments during the course of the evolution of a control structure was examined. This method was compared to the method of ‘fine tuning’ in a physical environment a control structure that has evolved in complete simulation. Interleaving physical and simulated environments improves performance of the eventual control structure. Possibly, this method allows for the evolved control structure to incorporate and retain advantageous behavioural patterns from both environments.

Keywords: evolutionary robotics, artificial evolution, real world evolution, simulation

1 Introduction

Evolutionary robotics is a field of science in which evolutionary methods are exploited in order to synthesise control structures for robots that operate in either a simulated environment or in the physical world. This approach is based on genetic algorithms, originally developed by [3]. In the presented study, a control structure for a physical robot was evolved. A number of epochs throughout the evolutionary process took place in simulation, while generations in between these epochs were evolved in a physical robot in a real environment. The issue under investigation is whether this method results in better performing control structures in comparison to the now more common method of evolving all generations in simulation entirely except for a small number of final generations.

1.1 Evolving control structures in simulated and physical environments

For reasons of efficiency control structures for physical robots are often partly evolved in simulation. The following method is often applied (see for example [6]): A control structure is evolved in simulation for a number of generations. Subsequently, the evolution-

ary process is continued in a physical environment for a smaller number of generations (typically about a tenth of the number of preceding generations). The advantage of such a method follows from the relatively small amount of time and work required for evolving in simulation when compared to real world evolution. [6] propose a number of techniques to optimise the results of the method. Unfortunately, these techniques are rather complicated and require quite some effort to improve the validity of the simulation.

Another way of combining simulation and physical robots when evolving control structure is by *interleaving* physical and simulated environments over the course of evolution. While assigning most of the work to the simulation, like the method mentioned above, an interleaving approach will *prepare* the eventual control system for the real world from the early stages on.

One notable example of the use of such an ‘interleaving’ strategy can be found in the work of [11]. In their experiment, the evolutionary process is split into separate phases, each of which handles a specific aspect of evolution. Their approach yields acceptable results, and is faster than if real robots had been used throughout [10]. However, it strongly deviates from Holland’s well proven GA method. In addition, human interference is required at every loop of the process, which may prove cumbersome.

1.2 Interleaving simulated and physical environments in evolution

Despite the problems found in the experiments of [11], the principle of interleaving simulated and real robots and environments may have a number of intrinsic pros when compared to, for example, the work of [6]. In the experiments of the latter, great efforts are made to model the physical robot and its environment as accurately as possible. In addition, a carefully selected type of noise must be added to the simulation. While these measures may actually optimise the results while minimising the number of real world trials that have to be run, they also demand for a certain amount of time, effort and expertise that may not always be available.

We propose a method that, like the method of [11] involves interleaving simulation and real world evolution during the course of evolution, but also allows for comparison with the method of [6]. The main difference with the method of [11] is that except for the nature of the environment and the robot (namely whether these are virtual or physical), all relevant factors are kept equal. Most importantly, the software that controls the robot and contains the control structure and the genetic algorithm and its parameters are identical. No efforts are made to optimally identify the simulation with the real world.

The main object of study in is the effect of interleaving on the fitness of eventual control structures. Over the course of evolution, control structures become increasingly reliant on the robot's environment and morphology. By breaking the chain of growing dependent on the simulation and by preparing, if you will, the control structure for the real world in bite-size chunks, a better performing robot is expected to result. One question that has to be answered first of all, is whether switching back and forth early on in the process improves eventual performance at all.

2 Experimental setup

2.1 The robot and the environment

The physical robots used in this experiment are composed of the parts found in the Lego Mindstorms robot kit. Its main component is the programmable *brick*, which provides motor outputs and sensory inputs, and in which the control structure is stored. A great advantage of Lego Mindstorms for this study is the availability of a customisable simulation plat-

form, the use of which is discussed in section 2.1.4. The descriptions below concern the physical robot and environment. The simulated robot and environment are modeled to their physical counterparts.

2.1.1 The environment

The robot was placed inside an arena in which it can freely move. This arena is surrounded by walls about the same height of the robot. Inside the arena, walls are placed as well. The robot is unable to cross any of these walls. It is possible however to drive around the walls *inside* the arena, since they do not form enclosed regions. The floor of the arena is white, except for the regions surrounding the walls, which are covered with strips of black paper. The white of the floor and the black of the paper constitute a difference in luminance large enough to be picked up by the light sensors of the robot. Light conditions were kept equal during the entire experiment.

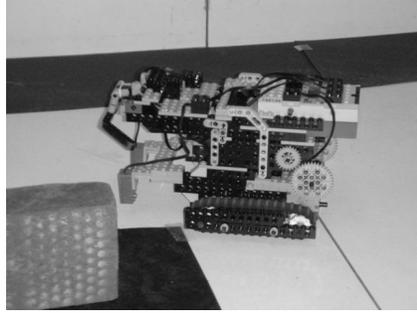
2.1.2 The robot

Figure 1(a) shows the morphology of the physical robot. Two motor units are connected to the Mindstorms brick, which control the two caterpillar tracks that are positioned at the sides of the vehicle. At the front side of the robot, three sensors are placed: two light sensors pointing to the floor and one touch sensor, which is activated when pressure is applied to a bumper in front of the robot.

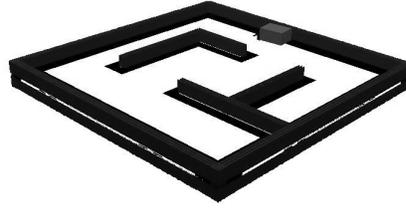
The program that was downloaded to the programmable brick contains the control structure for the robot, which processes incoming information from the sensors and generates output patterns which result in behaviour. Additionally, the software evaluates the robot's performance by monitoring the output patterns and keeping track of potential bump events. The details of the evaluation method are described in detail in §2.2.

2.1.3 The control structure

The control structure of the robot in this experiment consists of an artificial neural network (ANN) that maps sensory input to motor output values. The ANN consists of a single layer with all-to-all connections, which means that every unit in the network receives input from every unit in the network, including itself. Of course certain connections may be effectively disabled by having a weight value of zero.



(a) The physical robot positioned in the arena



(b) The simulated environment

Figure 1: The physical environment and the simulation

The weight values of the network and the biases of the units are subject of evolution.

The network contains 10 units. Therefore it can be represented by a 10×10 matrix in which the weights of the connections between the units are stored and an array of size 10 which holds the biases of the units. The values of the two light sensors are treated as input to two of the units of the network. The output values of two other units are used to set the power of the motors. The value of the touch sensor is not fed into the network, but used only for determining the genotype fitness (see §2.2.1).

Learning does not take place during the life time of the network, so the weight and bias values remain constant for each individual. At each time step t the activation of each unit i is updated by summing the products of the activation all incoming units j and the corresponding weights w_{ji} . Each unit has a bias $bias_i$ which is added to this sum before the result is passed through a sigmoid function which puts the activations on a curve between -0.5 and 0.5 .

Input values from the light sensors are treated as activations coming into the network and are thus added to the net input of the corresponding input units.

2.1.4 The simulation

The Lego Mindstorms Simulator package (LMS, [4, 9]) was used for modeling the robot and its environment and to simulate its behaviour. This simulation platform allows for detailed modeling of both the environment and the morphology of the the robot. In addition, control structures developed for LeJOS [1, 2], an operating system for Mindstorms robots that allows for the execution of Java programs, can

be used without modification in real robots as well as in robots simulated through LMS.

2.2 Evolutionary process

The evolutionary process applied to the control structure of the robot in this study is based on the ideas concerning genetic algorithms (GAs) as proposed by [3] and the work of [7] and [12] on the appliance of GAs to neural networks and that of [8] on Evolutionary Robotics.

The weights and biases to which evolution is applied were encoded into a genotype, constituted by an array of 110 floating point numbers (100 representing the weights and 10 for the bias values). A control structure can be extracted easily by instantiating a neural network with the values from the array filled in properly.

There were 20 genotypes in each generation. The first generation consisted entirely of randomly generated genotypes. Subsequent generations contained a copy of the genotypes of the ten best performing individuals of the previous generation, complemented by *mutated* versions of the same ten genotypes. Mutation is a genetic operator proposed by [3]. It is essential for the evolutionary process, since it yields minor variations in the genotype, which may result in improved performance in the resulting individual. Within the mutated genotypes, each number was selected for mutation with a chance of .05. If a gene was selected for mutation, its value was altered by a random number from a gaussian distribution with $\mu = 0$ and $\sigma = .3$.

2.2.1 Fitness function

Selection of the individuals to be kept for following generations took place on basis of the *fitness* of each genotype. The fitness of a genotype was determined by assessing the performance of the resulting control structure. This process can be described by the following steps:

1. Fitness is set to zero
2. At each time step while the robot is operating:
 - (a) If the robot is driving backwards, the square root of the product of the *power* values of both motors is *subtracted* from the fitness; otherwise, if no bumps have occurred in the past 100 cycles, the square root of the product of the *power* values of both motors is *added* to the fitness.
The *power* values of the motors are values that represent the speed at which the respective wheels are set to turn. Values of *power* < 0 indicate backward motion of the wheels. The *power* values are calculated from the activations of the output units, resulting in values between -50 and 50 .
 - (b) If the bumper is touched, a penalty value is subtracted from the fitness. The penalty is a constant measure we set to the value of 50 .
3. If $fitness < 0$: $fitness = 0$
4. Deliver fitness

The robot got to operate for a fixed number of time steps. In simulation mode, the number of steps was set to 500. The physical robot was programmed to report its fitness value after 300 time steps. This was done for reasons of efficiency. Since in this study it is not the intention to compare fitness values from the simulation and the physical environment, the difference is irrelevant.

2.2.2 Comparability of real world fitness and simulation fitness

It is not the aim of this study to directly compare fitness values of the physical environment and fitness values of the simulated environment. Therefore, no measures were taken to guarantee the comparability

of the fitnesses. Identically performing robots (assuming such a thing exists) in different environments (one in simulation, one in the real world) are unlikely to obtain even remotely comparable fitness values. An important cause for this difference, is the difference in the number of time steps between environments. Furthermore, there are numerous factors beyond verification of control that might influence the fitness measurements and are likely to differ between environments. Therefore, and since no intent is made to make comparisons between environments, one should consider ‘real world fitnesses’ and ‘simulation fitnesses’ to be on independent and incomparable scales.

2.3 Conditions

The experiment consisted of two conditions, in each of which a control structure was evolved over 55 generations, either fully in simulation or in an interleaved fashion, followed by five final generations of evolution in the real world. The results of the experiment are based on a comparison of the performance of the robot in the five final generations in both conditions. For a schematic overview of both conditions, see figure 2.

2.3.1 Complete simulation

In the first condition, the control structure was evolved in simulation over the first 55 generations. Five generations of real world evolution followed. This method is similar to the one used for example by [6].

Computation of the Pearson coefficient showed a reasonable correlation between generation and fitness for the first 55 generations ($r = .27$) as well as for the final five ($r = .24$), which indicates that the fitness has improved over the course of evolution, in both simulation and real world.

2.3.2 Interleaved evolution

In the second condition, epochs of evolution in simulation and evolution in the real world were alternated. The first ten generations were evolved in simulation, followed by a series of five generation evolved in the physical environment. This pattern was repeated until 60 subsequent generations had been evolved in total, of which 40 were the result of simulated evolution and 20 were evolved in the real world.

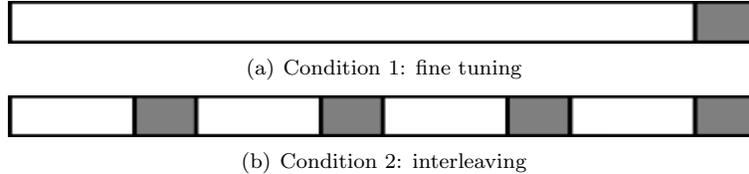


Figure 2: Schematic presentation of the two conditions of the experiment. The bars represent the entire course of evolution. The white zones stand for evolution in simulation; the grey parts represent evolution in the physical environment.

In this condition, reasonable correlations were found as well. For the entire set of generations evolved in simulation, a correlation of $r = .35$ between generation and fitness was found; the generations evolved in the physical environment yielded a correlation of $r = .20$.

2.4 Results

A robot control structure was evolved in two conditions. In the first condition, evolution in simulation and evolution in a physical robot were interleaved during the first 55 generations. In the second condition, the control structure was evolved in simulation entirely over the first 55 generations. To test whether interleaving simulation and real world evolution affects the performance of an eventual control structure for a physical robot, five final generations were evolved in a physical robot in both conditions.

Condition	N	Fitness	
		Mean	SD
Fine tune	100	142.3	260.6
Interleaved	100	1586.1	1067.0

Table 1: Descriptives

The fitnesses of the genotypes from these final generations were compared. As can be seen in Table 1, the mean fitness of the ‘interleaved’ condition over the final generations is about ten times as high as that of the ‘fine tune’ condition. In Figure 3(b) one can see that the average fitness of the ‘interleaved’ condition is in fact higher in each of the final generations. An analysis of variance over the final five generations showed that the observed difference between conditions is significant ($F(1, 198) = 172.76$, $p < .00$).

This means interleaving simulation and real world over the course of evolution positively affects the

performance of the eventual control structure. The strength of the effect is reasonable ($\eta^2 = .28$).

2.5 Conclusion

The results of the experiment described above show that interleaving simulation and a physical environment over the course of evolving a control structure for a physical robot positively affects the fitness of control structures subsequently evolved in a physical robot.

These results support the claim that a control structure in evolution can be prepared for its eventual habitat from the early stages of evolution on. However, it cannot be concluded that the results are completely due to the interleaving per se. One should consider the fact that the sheer number of generations evolved in the physical world (and thus in the environment in which the ultimate comparison was done) differs greatly between conditions. This can be seen clearly in Figure 2. The results from this study merely indicate the advantage of alternating environments during evolution in comparison to evolution through a long series of evolution in a single environment. What cause underlies this effect should be investigated in further research.

3 Discussion

3.1 Simulation/real world ratio

Evolving control structures in simulation requires radically less time and labour than doing so in a physical environment. This is the main reason for utilising simulation when evolving a control structure for a physical robot. Switching back and forth between simulation and the real world proves to be advantageous for the suitability of the ultimate control structure, but by definition requires quite some work to be done outside the simulation.

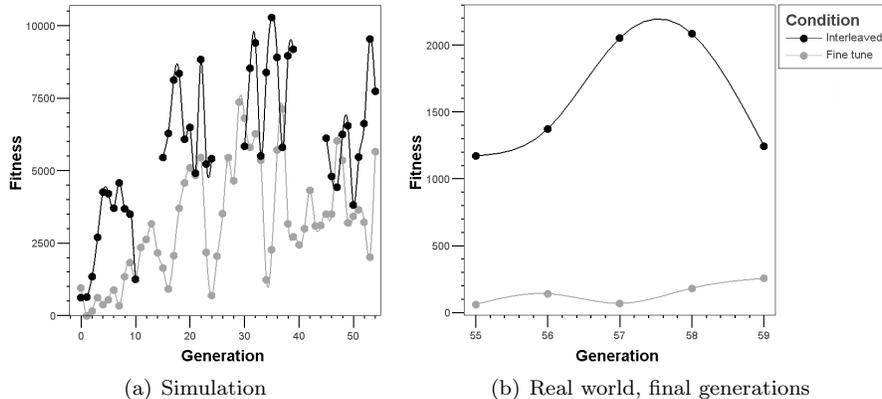


Figure 3: Fitnesses for both conditions. The dark line shows the mean value for each generation in the ‘interleaved’ condition; the grey line represents the ‘fine tune’ condition.

Efficiency of robot evolution can be defined on two scales. First, there is optimality of behaviour – or performance. Performance can be thought of in terms of fitness, or some other evaluation function of the behaviour of the (ultimate) control structure. Second, the amount of time and labour required over the course of the process are relevant. This factor can be called ‘costs’.

A control structure evolved entirely in the real world can be expected to behave near-optimally. However, costs are as high as they can be. The opposite is the case in simulation-only evolution. This method has very low costs, but yields poorly performing robots. Depending on one’s requirements, a certain balance between costs and optimality is usually desired. By altering the frequency and length of real world intermezzi the costs and performance can be fine-tuned. It would be interesting for further research to study the exact relation between these factors.

3.2 Scope of influence of a real world epoch and positioning effects

An issue related to the simulation/real world ratio is the scope of influence of one real world epoch: for how many generations of subsequent simulation trials are the genetic patterns that work out well in the real world retained? If a real world epoch early in the course of evolution would have a positive influence on performance in the final generations (like in the configuration sketched in Figure 4(b)), the effect of one such epoch could be concluded to be quite strong and long-lasting. If a significant effect would on the other

hand only be found where real world epochs follow each other closely (4(c)), the scope would be quite small.

The evolutionary process might also depend on the position in time at which alternations take place. Possibly the effect of an early real world epoch (Figure 4(b)) is stronger than that of a late one (4(c)), regardless of scope effects.

Both issues can be investigated in further research, in which one would have to carefully distinguish between positioning effects and effects of the scope of influence. These issues are closely related, but should not be confused. The findings of such investigations are of great relevance to the interleaving method of robot evolution; from their results, an optimal interleaving configuration could be constructed. In addition, these results could shed some light on the issue discussed in the following section.

3.3 General control structure or evolutionary memory?

From the results of the experiments in this study it shows that moving the evolutionary process from the simulation to the real world a number of times before finalising the process in the real world is beneficial. This effect implies that the improvements obtained during early real world epochs are somehow retained over the remaining course of evolution, including the stages in which evolution takes place in simulation. This phenomenon is particularly striking since many parameters that establish robot behaviour are expected to conflict with each other in the different environments. For example, specific assumptions about

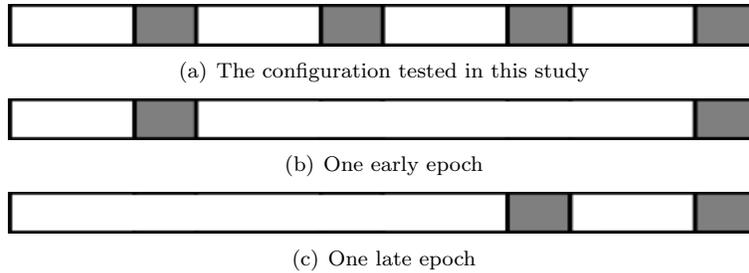


Figure 4: The scope of influence of one real world epoch

light conditions are particularly useful, but cannot be expected to hold in both environments.

The retention of beneficial genetic properties (and therefore: fitness) can be explained in two ways. One possibility is that the control structure simply grows more general: one can very well imagine a control structure that somehow allows for operation in either environment, conducting robot behaviour through very general methods. One example of a rather specific and a more general method can be given for the task of detecting a visual border in a black and white scene. In a specific environment, it may suffice to respond when light intensity drops below a certain threshold. This method would not work however for a wide range of environments, in which light conditions and perhaps even sensor properties could vary. A more general approach would be to attend relative changes in the observed intensity, and recognise borders by sudden changes in intensity. Even more general approaches can be thought of. It is possible that the generality of a control structure increases during evolution.

In practice however, increased generality is bound to reduce optimality for a specific situation. It is evident that there are no general methods for optimal behaviour in most natural or pseudo-natural (e.g., the arena of the experiment in this study) environments under changing conditions. Still, the control structure evolved in the experiment is not necessarily optimal for either environment, so generality could very well explain why fitness is not lost when the environment changes.

Another explanation can be found by drawing an analogy with nature. While habitats change over time, the genetic makeup of a species changes along. Morphological and behavioural properties of the phenotype gradually adapt to new dangers and opportunities provided by the environment. While many properties lose their function over time, it does not

mean they are genetically ‘overwritten’ immediately. Some archaic properties, such as androgenic hair still manifest themselves in modern human beings, without their function in modern life being clear. In addition, many properties might lay hidden in the genes that once had a distinctive function, but now no longer even manifest themselves.

The retention of archaic behavioural and morphological patterns may have proved to be an evolutionary advantage in cases of reoccurring environmental changes, such as ice ages. What principle could underlie such a retention mechanism? One possibility might lay in the vast portions of the DNA of many species that are not expressed in the corresponding phenotype. This so called ‘junk’ DNA constitutes up to 98% of mamallian genomes [5]. The function of ‘junk’ DNA in biology is under debate and it is not the aim of this thesis to contribute to this discussion in any way. Still in the experiments described earlier, parts of the genotype might, instead of or in addition to coding for the robot’s behaviour, implement something that could be called *evolutionary memory*: the storing of genetic patterns in such a way that they are accesible for incorporation in the phenotype in future times. This phenomenon can be explained as an emergent property of genetic mutation in somewhat oversized phenotypes; once a pattern has settled, the chances of it disappearing completely – even after it has become obsolete – are relatively small.

In my opinion, evolutionary memory can account for the retention of fitness as well. It would be very interesting to investigate whether evolutionary memory or increasing generality can best explain this phenomenon. An analysis of the development of the genotypes could shed some light on this issue. If evolutionary memory is exploited, certain patterns should be expected to arise during real world epochs and be retained during the remaining course. If such an effect cannot be found, the control structure is

probably becoming more general.

It would be very interesting to know whether something like evolutionary memory is exploited. This would be of great interest to the field of evolutionary robotics, and possibly to the study of evolution in general as well. Even though genetic computation as applied in this study is an extremely simplified model of biological evolution, effects that are found in it could reflect properties of actual evolution.

References

- [1] B. Bagnall. *Core LEGO MINDSTORMS Programming: Unleash the Power of the Java Platform*. Prentice Hall, Upper Saddle River, NJ, 2002.
- [2] G. Ferrari, A. Gombos, S. Hilmer, J. Stuber, M. Porter, J. Waldinger, and D. Laverde. *Programming LEGOMindstorms with Java*. Synpress Publishing, Rockland, MA, 2002.
- [3] J. H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA, 1975.
- [4] B. Künsting. Konzeption und prototypenhafte implementation einer lego-mindstorms simulationsumgebung - lejos-emulation. Master's thesis, Universität Paderborn, 2004.
- [5] W. Makalowski. Genomic scrap yard: how genomes utilize all that junk. *Gene*, 269:61–67, 2000.
- [6] O. Miglino, H. H. Lund, and S. Nolfi. Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4):417 – 434, 1995.
- [7] M. Mitchell. *An Introduction To Genetic Algorithms*. MIT Press, Cambridge, MA, 1996.
- [8] S. Nolfi and D. Floreano. *Evolutionary Robotics*. MIT Press, Cambridge, MA, 2000.
- [9] W. Sträter. Konzeption und prototypenhafte implementation einer lego-mindstorms simulationsumgebung - simulationsmaschine. Master's thesis, Universität Paderborn, 2004.
- [10] J. Walker, S. Garrett, and M. Wilson. Evolving controllers for real robots: A survey of the literature. *Adaptive Behavior*, 11(3):179 – 203, 2003.
- [11] M. S. Wilson, C. M. King, and J. E. Hunt. Evolving hierarchical robot behaviours. *Robotics and Autonomous Systems*, 22:215 – 230, 1997.
- [12] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87:1423 – 1447, 1999.