On the Influence of Phenotype Plasticity on Genotype Diversity

Ingo Paenke
Institute AIFB
University of Karlsruhe
D-76128 Karlsruhe, Germany
Email: paenke@aifb.uni-karlsruhe.de

Jürgen Branke
Institute AIFB
University of Karlsruhe
D-76128 Karlsruhe, Germany
Email: branke@aifb.uni-karlsruhe.de

Yaochu Jin
Honda Research Institute
Europe GmbH
D-63073 Offenbach/Main, Germany
Email: yaochu.jin@honda-ri.de

***Abstract*— A large body of research has investigated the advantages of combining phenotype adaptation and genotype adaptation. The hybridization of genetic search and local search methods, often known as memetic algorithms, and the influence of learning on evolution, i.e., the Baldwin effect and the Hiding effect, have been widely studied. However, most work assumes a stationary environment, and thus overlooks potentially advantages or disadvantages that can arise from phenotype plasticity only in changing environments. We show that a process with two levels of adaptation allows the system to operate on two different levels of diversity at the same time, which can be of great advantage under certain environmental conditions.**

I. Introduction

In nature, species need to cope with continual environmental changes. Evolution has found a variety of adaptation mechanisms that seem to be tailored for the particular environmental challenges, among which the individual-level and population-level adaptations are two main mechanisms [1]. While in individual-level adaptation, a single individual directly reacts to environmental changes, e.g., through *phenotype plasticity*, population-level adaptation requires a collective adaptation of a population usually by means of *genotype adaptation*, which is dependent on *genotype diversity*. For convenience, we will use genotype adaptation and phenotype plasticity in this paper to denote the population-level and individual-level adaptations.

The two levels of adaptation may interact with each other explicitly or implicitly. Two types of interactions, often known as *Baldwinian inheritance* and *Lamarckian inheritance*, have been discussed in artifical and biological evolutionary systems. In the Baldwinian inheritance, phenotype changes of a parent cannot be directly inherited by its offspring, while in Lamarckian inheritance, the changes on the phenotype are directly encoded back onto the genotype and can be passed on to offspring. It is widely believed that the Lamarckian inheritance is biologically less plausible, however, the Lamarckian inheritance can be used in an artificial evolutionary system if the mapping from the genotype to phenotype is surjective and invertible. Interestingly, it has been demonstrated within an artificial evolutionary systems that the Lamarckian inheritance mechanism performs better in stationary environments, whereas the Baldwinian inheritance mechanism works better in changing environments [2]. In this paper, we will investigate the Baldwinian inheritance mechanism, not only because of its biological plausibility and the findings in [2], but also because it does not restrict our discussions to evolutionary models in which the genotype to phenotype mapping is surjective and invertible.

Given the fact that for almost all species in nature adaptation takes place on both individual and population levels, one could ask if this principle is in general inevitable for adaptation to a changing environment. In the field of evolutionary computation, methods based on genotype-only adaptation have been used successfully both for solving stationary optimization problems, and dynamic ones, refer to [3]. Note, however, that a common feature of the evolutionary algorithms (EAs) for solving dynamic optimization problems is that an additional mechanism must be included to maintain a sufficient population diversity in order to avoid premature convergence.

In this paper, we revisit the question of what advantages a combined genotype-phenotype adaptation has compared to a genotype-only adaptation. In the literature, the following two aspects concerning the combination of genotype adaptation and phenotype plasticity have been discussed.

First, a combined genotype-phenotype adaptation allows to exploit distinct adaptation mechanisms at the genotype and the phenotype levels. In nature, evolution (gradually and slowly) changes the structural aspects of an organism while phenotype plasticity plays the role of fine-tuning the parameters for a given structure. Similarly in evolutionary optimization, a global evolutionary search (genotype adaptation) is often combined with a local learning algorithm (phenotype adaptation), which is usually intended to address the *exploration and exploitation dilemma* [4]. Evolutionary algorithms combined with a local search technique are often known as memetic algorithms [5].

Second, phenotype adaptation may change the evolutionary pathways, an effect that has first been formulated by Baldwin [6] and later termed the *Baldwin effect* [7]. The first computer simulation of the Baldwin effect demonstrated that the Baldwin effect can indeed guide evolution towards a global fitness optimum [8]. Recently, a mathematical framework, called the gain function has been developed that allows to predict whether phenotype plasticity accelerates or decelerates evolution [9]. The main idea of the gain function approach

is to determine whether phenotype plasticity enlarges fitness difference between fit and unfit individuals or reduces the difference. The latter effect has first been mentioned in the biology literature [10], and was later termed *Hiding effect* [11]. Actually, the Hiding effect is a special case of the Baldwin effect, and it cannot be said *per se* that it decelerates evolution. It is shown in the analysis in [9] that the hiding effect does slow down the evolution on a unimodal fitness landscape. Recently, it has been demonstrated empirically on a bi-modal fitness landscape [12], and mathematically based on random walk theory [13] that the Hiding effect may accelerate evolution on a multi-modal fitness landscape, since phenotype plasticity tends to smoothen fitness valleys.

The above advantages of the combined genotype-phenotype adaptation have been observed in models where a stationary environment is assumed. They hold in principle in dynamic environments as well. One aspect that is particularly relevant to changing environments has been demonstrated in [14]. In that work, the well-known model of Hinton and Nowlan [8] is extended to dynamic environments and a plasticity cost component is also included. Based on simulation and mathematical results, it is concluded that phenotype plasticity evolves if the plasticity cost is low.

This paper demonstrates another effect of combining phenotype plasticity with genotype adaptation. It is found that phenotype plasticity enables the evolutionary system to operate with two different levels of diversity simultaneously. In Section II, we introduce a simulation model that takes the cost of phenotype plasticity implicitly into account. In particular, we assume a trade-off between generation turnover, i.e., the average number of offsprings per time-unit, and an individual's lifetime. The model uses similar adaptation mechanisms for genetic adaptation and phenotype plasticity in order to avoid a serious bias toward a particular adaptation mechanism. A detailed analysis on population diversity is presented in Section III to investigate the influence of phenotype plasticity on genotype diversity. Based on the insights gained in the diversity analysis, we study the adaptation behavior of the genotype-phenotype adaptation and the genotype-only adaptation in changing environments in Section IV. Conclusions are given in Section V.

II. Simulation Model and Examples

In Fig.1 one time step (t) of the discrete-time evolution model used in this work is illustrated. A population P_t is composed of n individuals, each having the same pre-specified constant lifetime L, a genotype value x, a (variable) phenotype value z, and correspondingly an adaptive value a. The adaptive value is the basis for phenotype and genotype adaptation defined by a time-dependent environment function f, i.e. $a = f(z,t)$. In evolutionary computation, this is often called *fitness function*. We do not use this term here to avoid inconsistencies with the biological definition of fitness, where fitness is considered to be proportional to the number of offspring, which can thus only be measured posthumously.

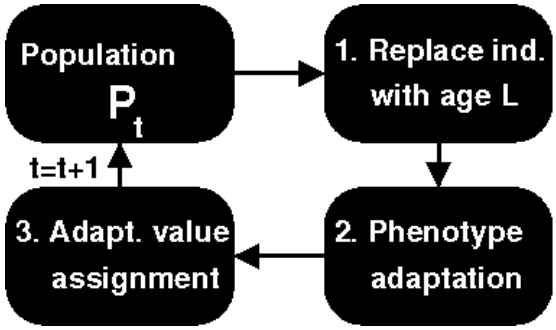

Fig. 1. The three main phases of the simulation model.

In each time-step, the following three phases can be distinguished (cf. Fig. 1):

Phase 1 - Replace individuals of an age L

All (m) individuals which have achieved an age of L, i.e., the end of their life, are removed from the population. Then m individuals are picked out as parents (probability proportional to the current adaptive value) to produce m offsprings, i.e., the genotype of the parents are copied to the offsprings. The genotype of the newly generated offsprings is then mutated by adding a random number sampled from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = \sigma_G$. The union of the original population and new offsprings without the perished individuals forms the new population.

Phase 2 - Phenotype adaptation

New-born individuals develop their phenotype first. This is done simply by applying the identity function to the genotype value, i.e., the innate phenotype equals the genotype value $z(x) = x$. All other individuals (of an age larger than or equal to 1) try to increase the adaptive value using a simple learning algorithm based on a rudimentary lifetime memory that stores the best phenotype found so far and the corresponding adaptive value that has been achieved at the time when the phenotype value was explored. When an individual randomly explores the environment, its phenotype is changed by adding a random number sampled from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = \sigma_P$. If the resulting new phenotype has a higher adaptive value than the memorized one, the individual adopts the new phenotype. Otherwise, the current phenotype is kept. The implementation of this kind of phenotype plasticity can be seen as a $(1+1)$-*EA* and is very similar to the genotype adaptation mechanism since in all simulations we set $\sigma_G = \sigma_P$, which avoids a strong methodological bias towards a particular adaptation mechanism.

Phase 3 - Adaptive value assignment

Finally, the individuals are assigned an adaptive value which depends on the environmental dynamics. Note that the adaptive value an individual has memorized for a certain phenotype value in Phase 2 may have changed due to environmental

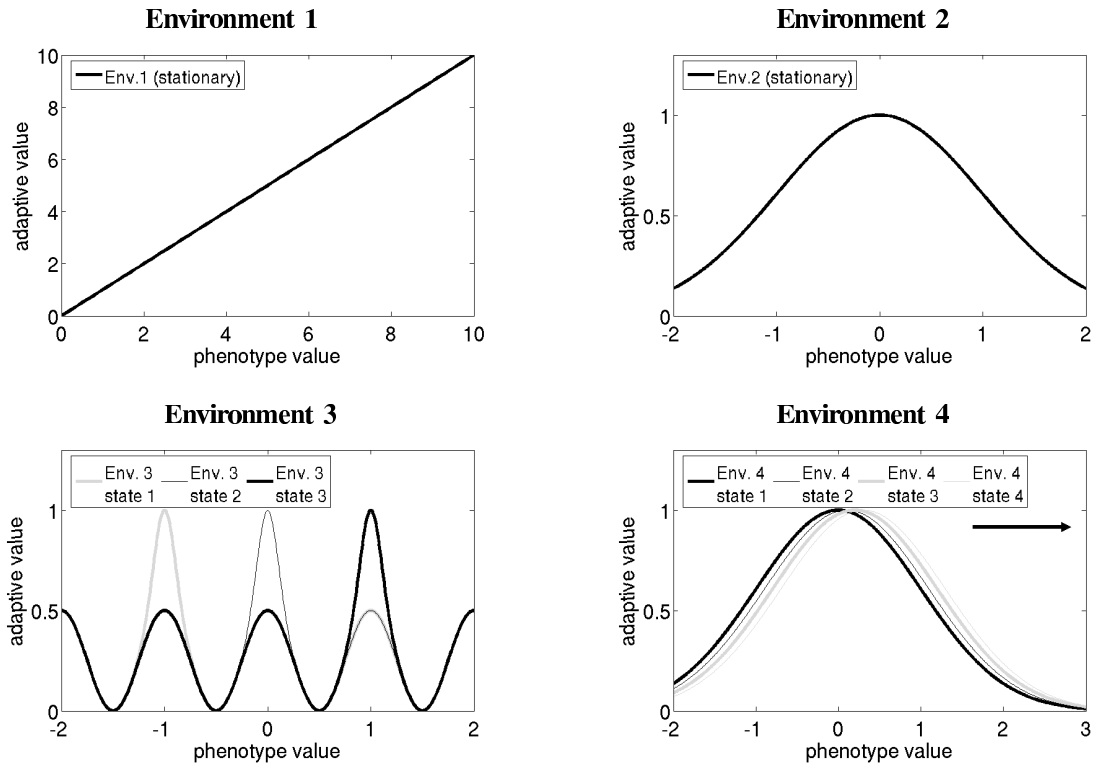

Fig. 2. Simulation environments of this paper, defined as the mapping from phenotype value to adaptive value. Env.1 and Env.2 are stationary, Env.3 and Env.4 are dynamic. There exist 3 states in Env.3, which change according to the pattern *state1-state2-state3-state2-state1-*... and w.r.t. to a change interval (cf. Section IV for a definition). There exists an infinite number of states in Env.4, where only the first 4 states are shown. The function continues to shift right-wards, with a velocity depending on the change interval.

changes. In Phase 1, the selection thereafter is based on the new adaptive value.

Model settings and main effects

This model naturally implements a trade-off between genotype and phenotype adaptation efforts. By keeping the overall adaptation effort constant and adjusting the lifetime L, we can distribute the adaptation efforts between genotype and phenotype levels. For example, with a constant population size 100, a lifetime $L = 10$ implies the replacement of 10 individuals at each time-step, whereas a lifetime $L = 1$ implies the replacement of 100 individuals at each time-step, which is identical to a conventional generational EA. In all simulations of this paper the population size is set to 100.

Within the model framework, the effects of including a phenotype adaptation (at the expense of genotype adaptation) are investigated by comparing the evolution of a population with a lifetime $L = 1$ to another one with $L = 20$. For the sake of readability, we simply denote the first case as *GO* (genotype-only) and the second as *GP* (genotype-phenotype) adaptation.

Changing the lifetime L (switching from GO to GP adaptation) has two obvious effects: First, since in the proposed model the population size is assumed constant, an increase in the average lifetime reduces the generation turnover rate. Secondly, a change in the average lifetime influences the phenotype adaptation, which in turn influences which individuals produce offsprings. In other words, phenotype adaptation causes a change in the mapping from the genotype to the adaptive value.

We now take a look at two simple stationary environments (f does not depend on t) which we denote as *Env.1* and *Env.2*.

Monotonically Increasing Environment (Env.1)

In Env.1 f is defined as a monotonically increasing linear function of phenotype z, in particular $f(z,t) = z$, cf. Fig. 2. The evolutionary goal is to "climb up" the (infinitely) monotonically increasing function as quickly as possible. Fig. 3 shows how in a typical simulation run, the genotype and the phenotype distribution change over time in cased of GO (left column) and GP (center and right column). Initially, the population is uniformly distributed on $[0; 1]$ and adaptation parameters are set to $\sigma_G = \sigma_P = 0.005$. The population with GO adaptation performs much better in this task. After 500 time steps, the mean population genotype value is beyond 12 while in case of GP adaptation, this value is below 4. This is caused by the reduction of generation turnover rate as described earlier. Clearly, in Env.1 phenotype adaptation is detrimental.

Stationary Gaussian Environment (Env.2)

In Env.2 fitness landscape f is defined as a Gaussian function centered at 0. The evolutionary goal is the convergence to the optimum (0), cf. Fig. 2. Adaptation

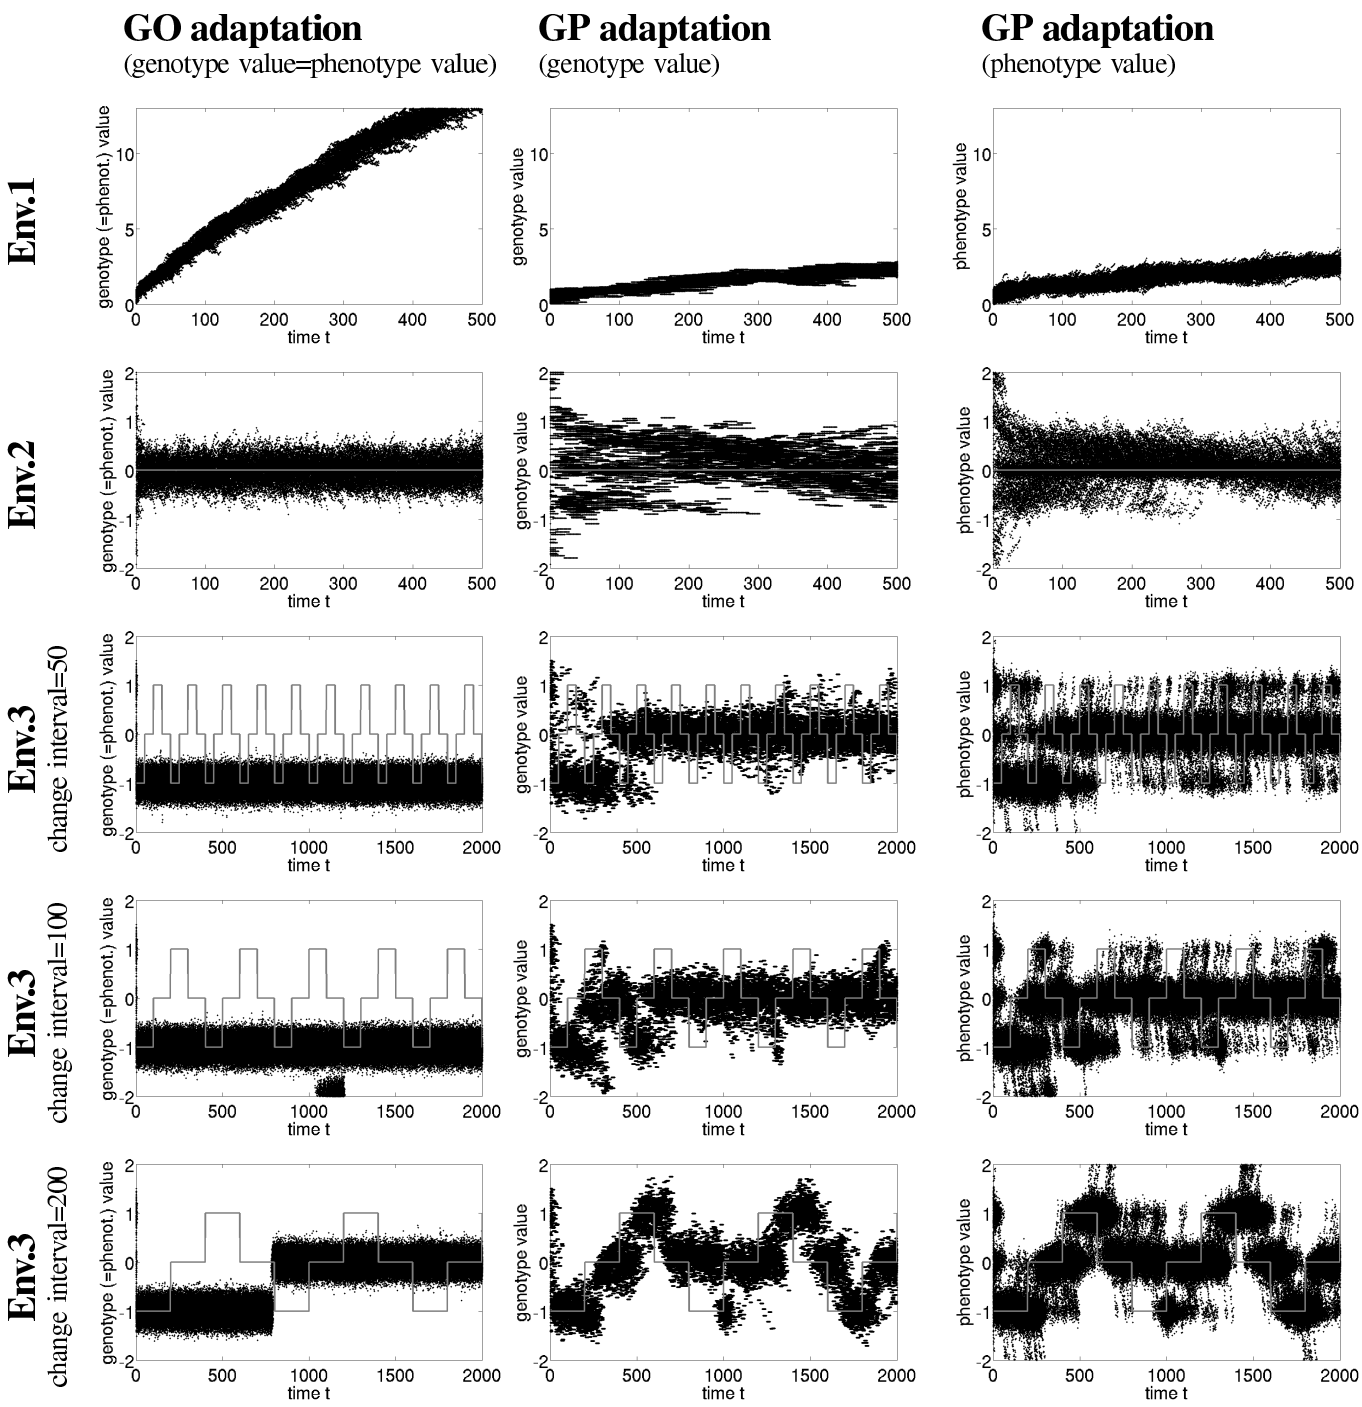

Fig. 3. Typical simulation runs for the three environments Env.1, Env.2, Env.3; Env.3 for different settings of the change interval, i.e., 50,100 and 200 (one row of panels for each setting, cf. Section IV for a definition of the change interval). The plots show the population distribution over time. All panels in the left column show the population distribution in case of GO (genotype-only) adaption, genotype equals phenotype in this case. All panels in the middle column show the genotype distribution in case of GP (genotype-phenotype) adaption and all panels in the right column show the phenotype distribution in case of GP adaption. The gray line marks the optimum (in case of Env.2 and Env.3).

parameters are set to $\sigma_G = \sigma_P = 0.005$. The population is initially distributed uniformly on $[-3; 3]$, thus several individuals are located in the proximity of the optimum already in the first generation. Fig. 3 shows typical simulation runs for Env.2. In case of GO, the population converges to a stable state after approximately 7 generations, while this takes much longer in case of GP. In both cases of GO and GP, full convergence is not achieved, although the population has reached a stable state. This phenomenon has been described as quasi-species [15]. This is the formation of a genotype or phenotype cloud around a wild type under mutation-selection balance. In this example, the wild type is $x = 0$ respectively $z = 0$.

It seems that the population in GO adaptation has a lower diversity w.r.t. both genotype and phenotype compared to GP adaptation case. Since diversity is a major requirement for adaptation in changing environments, we continue with a detailed diversity analysis of the simulation data of Env.2.

III. Diversity Analysis

Diversity is commonly defined as *the variety and abundance of organisms at a given place and time* [16]. It is accepted wisdom that a certain level of diversity is a condition for adaptation to new environmental conditions, however, it is less clear how to measure it [16], [17].

Measurements

Numerous diversity indices have been suggested in the biology literature, but it was shown that no single index is universally superior [18]. In computational intelligence, diversity has been studied too, mainly in artificial life and evolutionary computation [19], [20]. Although developed for different purposes, all diversity indices reflect at least one of the two aspects, namely, richness and evenness. By richness, it is meant how many different elements exist in an element set, e.g., how many species in the ecosystem or how many different genotypes or phenotypes in a species or a local population. Evenness refers to the distribution of a given set of elements w.r.t. certain element properties. A uniform distribution produces maximum evenness. Some diversity indices focus on either of the two properties, e.g. the Hurlberts Probability of Interspecific Encounter [21] quantifies evenness. Most indices aggregate both aspects, such as the Shannon entropy [22], [23] and the Simpson's index [24]. Other indices are developed in discrete domains and no counterpart in continuous spaces exists. If a continuous space is involved, as in our model, these indices can still be used by discretizing the continuous space. However, the choice of the discretization parameters can bias the measurement. In order to cope with such biases, we will adopt two indices in this work. The first index we use is the well-known Simpson index [24], which is widely used in biology and has been recommended in [18]. The Simpson index H_1 reflects the probability that two randomly sampled individuals are not equal (do not belong to the same partition class) and is defined as

$$H_1 = 1 - \sum_{i=1}^{m} \frac{n_i}{n} \,, \tag{1}$$

where n is the population size, and n_i the number of individuals in partition class i (out of m partition classes). H_1 increases with both evenness and richness, which is a desired property.

The second index (H_2) that we adopt here is parameter-free and is defined directly on a continuous space as follows:

$$H_2 = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1}^{n} |x_i - x_j| \,, \tag{2}$$

where x_i is the genotype (or phenotype) value of individual i. H_2 measures the average Euclidean distance of individuals within a population. Using average distance as diversity measure is also typical in the field of evolutionary computation with a binary representation, where the average Hamming distance between the genotypes is often used for measuring diversity. It should be pointed out that counter examples can be constructed in which a higher evenness or a higher richness does not increase the value H_2.

Fig. 4 shows diversity measurements (H_1 and H_2) over simulation time for Env.2. Note that for the Simpson index (H_1) we have discretized the genotype and phenotype spaces into partition classes $(-\infty; -3], (-3; -2.75], (-2.75, 2.5], \ldots, (2.75; 3], (3; +\infty)$.

Convergence process diversity

In Env.2, the genotype diversity decreases more slowly with GP than with GO. Several factors cause this behavior.

First, the reduced generation turnover of GP (compared to GO) obviously slows down the diversity decrease. This can be seen by comparing the slope of the thick grey line (GO) with the slope of the thin dashed line, which represents evolution with individuals of a lifetime $L = 20$, where the phenotype adaptation during the 20 life time units is disabled, hence avoiding the Hiding effect. Clearly, the thin dashed line decreases more slowly.

However, there is a second aspect, which is related to the selection-reproduction pattern of GP adaptation. In particular, the algorithm that represents the GO case is known as (conventional) generational EA whereas the algorithm that generated the thin dashed line ($L = 20$, no phenotype adaptation) is known as a steady-state EA (note, however, that the oldest parents are replaced here). For these two algorithms it has been shown that the generational EA produces a larger degree of diversity than the steady state EA even if both EAs use the same number of fitness evaluations, where entropy has been used as the diversity measure [25]. Our simulation results support these findings: In Fig. 4 (bottom-right panel) we plot every 20 time steps of the dashed line of the top-right panel of Fig. 4, thus showing the effect of 100 offspring replacements (for 20 time steps, in each time step the 5 oldest individuals are replaced, $20 \cdot 5 = 100$). We compare this with the thin black

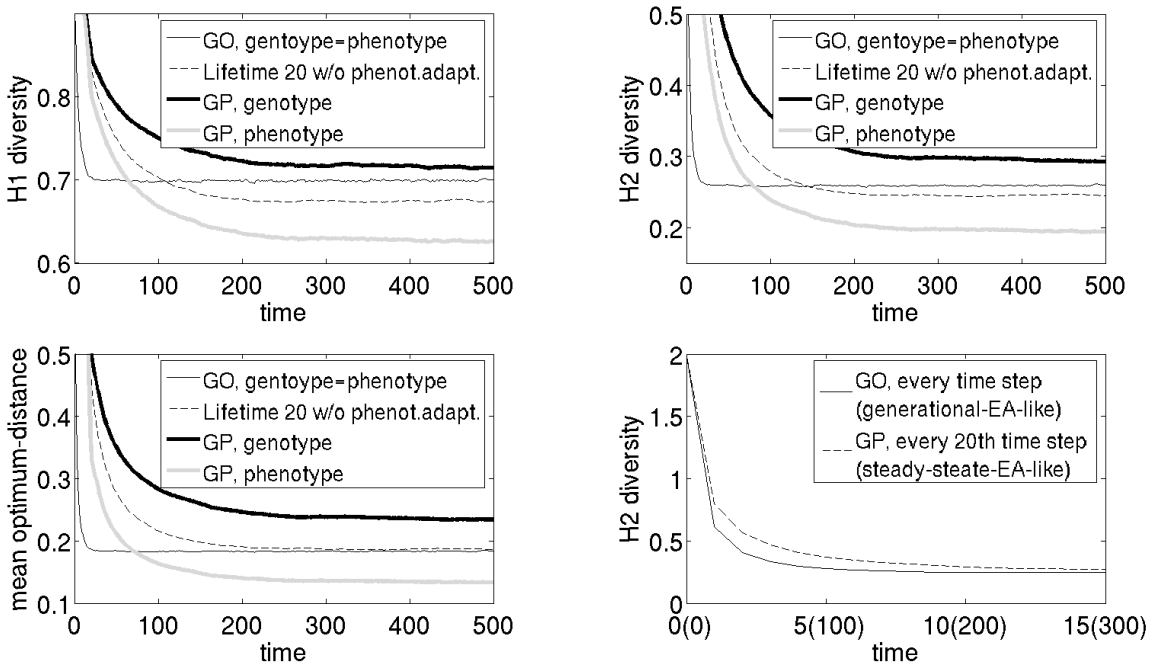

Fig. 4. Averaged diversity data of 1000 simulation runs in Env.2. **Top-left:** H_1 diversity, **Top-right:** H_2 diversity, **Bottom-left:** Mean individual distance to the optimum (0), **Bottom-righ:** H_2 comparison of generational-EA-like and steady-state-EA-like evolution (see text for details). Graphs refer to the following. **Black thin solid line:** Genotype/phenotype in case of GO (genotype-only) adaption evolution i.e. (genotype and phenotype are equal in this case), **Black thin dashed line:** Genotype/phenotype of individuals with lifetime $L = 20$ with disabled phenotype adaption, i.e., phenotype is constant throughout the 20 lifetime units, **Black thick solid line:** Genotype in case of GP (genotype-phenotype) adaptation, **Gray thick solid line:** Phenotype in case of GP (genotype-phenotype) adaptation.

curve of top-right panel of Fig. 4 (where one x-axis unit also shows the effect of 100 replacements). As expected, Fig. 4 (bottom-right panel) shows that the diversity of the steady-state EA is lower than that of the generational EA before and after convergence.

The third factor is related to the change in the mapping from the genotype to the adaptive value introduced by the phenotype plasticity. This effect can be seen by comparing the thin dashed line with the thick black line (genotype in case of GP) in the H_1 and H_2 graphs (Fig. 4, top-left and top-right panels). It can clearly be seen that enabling phenotype adaptation slows down the genotype convergence. This must be attributed to the Hiding effect that we discussed in Section I, i.e., a decelerated evolutionary process caused by a reduction of selection pressure resulting from the phenotype adaptation. Note that in principle, phenotype adaptation can also increase selection pressure. We refer to [9] for details, where exact mathematical conditions for both cases have been derived. In conclusion, the following observations can be made:

1) *Phenotype adaptation slows down the loss of genotype diversity due to reduced generation turnover rate.*
2) *Phenotype adaptation slows down the loss of genotype diversity due to a steady-state-EA-like generation overlap.*
3) *Phenotype adaptation slows down the loss of genotype diversity in (the likely) case of the Hiding effect.*

Post-convergence diversity

By simulation time 300 all curves remain more or less constant, which allows us to identify the effect of a combined genotype-phenotype adaptation on the *post-convergence* quasi-species diversity.

We can see that the genotype diversity is higher in case of GP than in case of GO, and at the same time the phenotype diversity is lower in case of GP. The reduced generation turnover is no longer an argument for this phenomenon. The reason for the rather high genotype diversity in case of GP is given by the occurrence of the Hiding effect. The influence of the Hiding effect on genotype diversity after convergence can exactly be seen by comparing the thin dashed line (for explanation see above) and the thick black line in the H_1 and H_2 graphs in Fig. 4. Clearly, the thick line remains at a higher level after convergence.

However, the phenotype diversity in case of GP is much lower. The explanation for this is straightforward, since the phenotype is (naturally) better adapted than the genotype. Additional evidence for this argument can be seen from the right-most panel of Fig. 4. In this figure, the mean distance to the optimum (0) is given, from which one can clearly see that the phenotype is on average closer to the optimum than the genotype.

In conclusion, the following observations can be made:

1) *Phenotype adaptation increases the (post-convergence)*

genotype diversity in (the likely) case of the Hiding effect.

2) *Phenotype adaptation helps the phenotype to settle closer to the optimum.*

However, no definite conclusion can be drawn on whether the inclusion of phenotype plasticity is positive or negative to evolution, as we will see in the following section.

IV. Changing Environments

In the preceding section we have shown how phenotype adaptation influences population diversity. Based on these results, we now demonstrate under what environmental conditions this influence is likely to have a positive and under what conditions this is likely to have a negative effect on the overall adaptation success.

Multi-modal dynamic environment (Env.3)

In Fig. 2 the environment dynamics of our first example of changing environments (*Env.3*) is shown. In this environment there are several local optima and one global optimum, which changes over time among three locations $\{-1, 0, 1\}$. Simulations have been conducted with different parameter settings. It turned out that (besides the lifetime setting, which is the same as in previous sections) two parameters, namely the *change interval* and the *adaptation step-size* have a strong impact on the overall adaptation process.

The *change interval* defines after how many time steps an environmental change happens. With a very small change interval, basically no adaptation takes place. In this case, the adaptation behavior is similar to that in an environment with many equally high optima. In particular, the population, which is initially distributed on many local optima, gradually disappears from the local optima due to genetic drift, and converges to a single local optimum. A large change interval causes a similar behavior to that in a stationary environment, except that after a (later) environmental change the population distribution is biased by the previous environmental state. We therefore present some intermediate settings of the change interval, namely, 50,100,200.

The *adaptation step-size* includes two parameters actually, namely the mutation step-size σ_G and the phenotype adaptation step-size σ_P. However, as argued earlier, we set $\sigma_G = \sigma_P$. With a very small adaptation step-size, adaptation is only possible through selection and is strongly dependent on the initial population distribution. With a very large adaptation step-size, adaptation is basically characterized by drift. After evaluating a large parameter set, the setting $\sigma_G = \sigma_P = 0.01$ turned out to be an interesting case for further investigation.

Fig. 3 (panels in rows 3,4 and 5) shows typical results for these settings. For all change interval settings of GO adaptation (panels in the left column) the population converges to a local optimum and is able to escape from it in only two situations. However the population is not able to move towards the global optimum. In case of a combined genotype-phenotype adaptation, the population manages to follow the optimum. In case of a quickly changing environment with a change interval of 50 (Fig. 3, Env.3, change interval=50), the steady state of the environment seems to be too short for the population to fully adapt to the global optimum. In case of a change interval of 100 (Fig. 3, Env.3, change interval=100), the steady state seems to be long enough for the population to adapt to the global optimum occasionally. Finally, with a slowly changing environment with a change interval of 200 (Fig. 3, Env.3, change interval=200), the population always follows the changing global optimum. Fig. 5 compares the average adaptive value between the GO and GP evolution for different change intervals, which provides further evidence that in environments like Env.3, the combination of genotype and phenotype adaptation brings about an advantage. While in the GO adaptation the population only by chance *matches* the global optimum (when it returns), the population always *adapts* to a changing optimum in GP evolution.

Environment with monotonic dynamics (Env.4)

In contrast to Env.3, the next example (*Env.4*) aims to demonstrate that combining genotype adaptation with phenotype adaptation can be detrimental to the overall adaptation process. The mapping from phenotype to adaptive value of Env.4 is described by a Gaussian function, which shifts periodically 0.1 units in positive direction as specified by the change interval (cf. Fig. 2). Fig. 5 shows the average adaptive value between GO and GP. In case of a quick environmental change, with a change interval of 1, both GO and GP fail to follow the optimum. At the other extreme, where the environmental change is slow, a change interval of 50, the population is able to follow the global optimum in both cases of GO and GP. However, at an intermediate level, with a change interval of 10, adaptation is successful in the case of GO but fails in the case of GP. This result can be attributed to the selection pressure reduction caused by phenotype adaptation. We conclude that the inclusion of phenotype adaption is detrimental if the environmental dynamics are monotonic, thus demonstrating that there exist environments in which GO adaptation is more appropriate.

V. Conclusion

We have shown that including phenotype plasticity in the evolutionary adaptation process influences the genotype diversity under the Baldwinian inheritance regime, when the overall adaptation "effort" is constant and similar adaptation mechanisms are used for genotype and phenotype adaptation. This is caused by the reduction in generation turnover and the change in the fitness landscape, which also influences the convergence process and the post-convergence behavior. We have identified two characteristics of dynamic environments that favor one or the other (genotype-only or combined genotype-phenotype adaptation). The advantage of the combined genotype-phenotype adaptation not only lies in its influence on the genotype diversity that opens an opportunity for long-term adaptation, but also results in a higher short-term adaptability, since the phenotype can be well adapted without a loss of the genotype diversity.

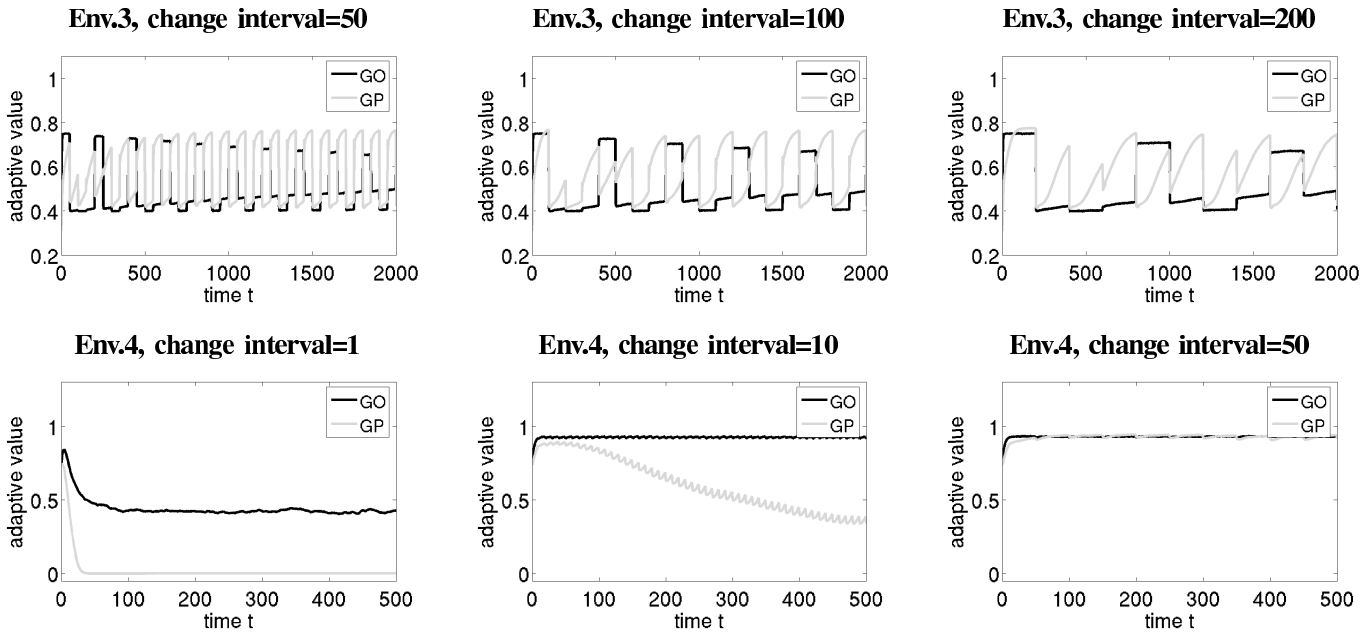

Fig. 5. Adaptive value for different change rates in Env.3 and Env.4, averaged over 100 simulation runs.

Our results may inspire new ideas in the field of artificial evolutionary systems where both a high level of diversity and a strong adaptation ability are important. In addition, we can also gain some insights into the evolution of biological systems in that our results shed some light on the question why a large variety of adaptation mechanisms exists among species.

In the future, we hope to show that the level of phenotype adaptability, i.e., the lifetime L in our model, can be evolved in a second-order adaptation process to fit well with a given dynamic environment.

ACKNOWLEDGMENT

The authors want to thank B. Sendhoff, E. Körner and H. Schmeck for their support.

REFERENCES

[1] L. Ancel and J. Bull, "Fighting change with change: Adaptive variation in an uncertain world," *Trends in Ecology and Evolution*, vol. 17, no. 12, pp. 551–557, 2002.

[2] T. Sasaki and M. Tokoro, "Comparison between Lamarckian and Darwinian evolution on a model using neural networks and genetic algorithms," *Knowledge and Information Systems*, vol. 2, no. 2, pp. 201–222, 2000.

[3] J. Branke, *Evolutionary Optimization in Dynamic Environments*. Kluwer, 2001.

[4] J. Holland, *Adaptation in natural and artificial systems*, 2nd ed. MIT Press, 1992.

[5] W. Hart, E. William, N. Krasnogor, and J. Smith, Eds., *Recent Advances in Memetic Algorithms*. Springer, 2005.

[6] J. Baldwin, "A new factor in evolution," *American Naturalist*, vol. 30, pp. 441–451, 1896.

[7] G. Simpson, "The Baldwin effect," *Evolution*, vol. 7, pp. 110–117, 1953.

[8] G. Hinton and S. Nowlan, "How learning can guide evolution," *Complex Systems*, vol. 1, pp. 495–502, 1987.

[9] I. Paenke, T. Kawecki, and B. Sendhoff, "On the influence of lifetime learning on selection pressure," in *Proceedings of Artificial Life X*, L. Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, and A. Vespignani, Eds., 2006, pp. 500–506.

[10] T. Johnston, "Selective costs and benefits in the evolution of learning," *Advances in the Study of Behavior*, vol. 12, pp. 65–106, 1982.

[11] G. Mayley, "Guiding or hiding: Explorations into the effects of learning on the rate of evolution." in *Fourth European Conference on Artificial Life*, P. Husbands and I. Harvey, Eds., 1997, pp. 135–144.

[12] R. Mills and R. Watson, "On crossing fitness valleys with the Baldwin effect," in *Proceedings of Artificial Life X*, L. Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, and A. Vespignani, Eds., 2006, pp. 493–499.

[13] E. Borenstein, I. Meilijson, and E. Ruppin, "The effect of phenotypic plasticity on evolution in multipeaked fitness landscapes," *Journal of Evolutionary Biology*, vol. 19, no. 5, pp. 1555–1570, 2006.

[14] H. Dopazo, M. Gordon, R. Perazzo, and S. Rissau, "A model for the emergence of adaptive subsystems," *Bulletin of Mathematical Biology*, vol. 65, pp. 27–56, 2003.

[15] M. Eigen, "Selforganization of matter and the evolution of biological macromolecules," *Naturwissenschaften*, vol. 58, pp. 465–523, 1971.

[16] A. Magurran, "Biological diversity," *Current Biology*, vol. 15, no. 4, pp. 116–118, 2005.

[17] S. Cousins, "Species diversity measurement: Choosing the right index." *Trends in Ecology and Evolution*, vol. 6, no. 6, pp. 190–192, 1991.

[18] R. Routledge, "Diversity indices: Which ones are admissible?" *Journal of Theoretical Biology*, vol. 76, no. 4, pp. 503–515, 1979.

[19] R. Morrison and K. D. Jong, "Measurement of population diversity," in *Selected Papers from the 5th European Conference on Artificial Evolution*, P. Collet, C. Fonlupt, J. Hao, E. Lutton, and M. Schoenauer, Eds., 2001, vol. 2310, pp. 31–41.

[20] C. Mattiussi, M. Waibel, and D. Floreano, "Measures of Diversity for Populations and Distances Between Individuals with Highly Reorganizable Genomes," *Evolutionary Computation*, vol. 12, no. 4, pp. 495–515, 2004.

[21] S. Hurlbert, "The nonconcept of species diversity: A critique and alternative parameters," *Ecology*, vol. 52, no. 4, pp. 577–586, 1971.

[22] C. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, 1948.

[23] E. Pielou, "Shannon's formula as a measure of specific diversity: Its use and misuse," *American Naturalist*, vol. 100, no. 914, pp. 463–465, 1966.

[24] E. Simpson, "Measurement of diversity," *Nature*, vol. 163, p. 688, 1949.

[25] F. Vavak and T. Fogarty, "A comparative study of steady state and generational genetic algorithms for use in nonstationary environments," in *AISB workshop on Evolutionary Computing*, T. Fogarty, Ed., no. 1143, 1996, pp. 297–304.