

Initial Modelling of the Alternative Phenotypes Hypothesis

Simon Worgan¹ and Rob Mills²

¹Information: Signals, Images, Systems Research Group, University of Southampton, Southampton, UK, SO17 1BJ

²Science and Engineering of Natural Systems Research Group, University of Southampton, Southampton, UK, SO17 1BJ
mmm05r@ecs.soton.ac.uk

Abstract

The alternative phenotype hypothesis contends that multiple phenotypes exist in a single genotype and are expressed by environmental or genetic cues. It further states that these multiple phenotypes will be maintained and improved in a population where the environment is unstable, in spite of the increased cost of this plasticity. In this work we propose a simple computational model to investigate the conditions under which alternative phenotypes become beneficial, and persist over evolutionary timescales. We find that the environment must vary to realise this hypothesis, and that these adaptations not only provide a fitness benefit in highly unstable environments but also continue to arise despite increasing stability and a corresponding gradual decline in fitness.

Introduction

The Alternative Phenotypes Hypothesis (APH), put forward by West-Eberhard (1986, 1989, 2003) puts across the view that phenotypic plasticity in the form of condition-sensitive phenotype expression (*i.e.* alternative phenotypes) is key in a sequence of evolutionary processes that lead to organic novelty, and in turn speciation and higher macroevolutionary events. Although the name may suggest that the evolution of stable alternative phenotypes is most of the operation, the hypothesis in fact covers several events of evolutionary significance.

Key to this hypothesis are alternative phenotypes, which are defined by West-Eberhard as: 'different traits expressed in the same life stage and population, more frequently expressed than traits considered anomalies or mutations, and not simultaneously expressed in the same individual' (2003, p377). In essence alternative phenotypes are when individuals from a single population can develop into different discrete phenotypes, and when this is environmentally cued it is also termed polyphenism. A familiar example is the dimorphism in sexually reproducing species (Lande, 1980).

In brief, the APH suggests that these alternative phenotypes arise when novel traits become established within a population. So long as each alternative phenotype is expressed in an advantageous environment it will be 'buffered' from negative selection; this can be thought of as akin to

diversity maintenance techniques employed in evolutionary computing. Potentially, over time each alternative phenotype becomes increasingly distinct, and if modified or localised conditions only favour one particular phenotype this could lead to the emergence of new lineages.

West-Eberhard's hypothesis has many stages and the supporting evidence provided is stronger for some portions than others. One aspect that is not as well substantiated is the specific conditions that afford polyphenic populations a selective advantage.

The APH is based on (extensive) surveys of experimental evidence connected together by verbal arguments. A complex multi-stage theory is sometimes hard to rigorously assess experimentally. Evolutionary simulation models of the processes involved in this hypothesis can assist in constructing more elaborate thought experiments, validating the consistency between stages of the argument, and help identify underlying mechanisms (Barandiaran and Moreno, 2006; Dennett, 1994). In this study we make some early steps towards these goals by modelling one stage of APH; specifically, the fixation of novel traits as viable alternative phenotypes.

The study of organism development is a large and active area of research in evolutionary biology (Wolpert, 2007; Hall, 1998), and is frequently the subject of studies in Artificial Life, *e.g.* Lindenmayer systems (Hornby and Pollack, 2001), ontogeny (Gear and Wiles, 2005), learning (Nolfi and Parisi, 1998) and phenotypic plasticity (Mills and Watson, 2006). However, much of this work focuses on aspects of evolvability arising from developmental representations. Rather than looking at the robustness of specific developmental trajectories themselves, this study focuses on a different aspect of development: the adaptive consequences of environmental influences on the development of several possible alternative phenotypes.

Accordingly, in this paper we propose a computational model to investigate when the evolution of stable, alternative phenotypes provides an adaptive advantage as compared to an evolving population that had no mechanism to support multiple phenotypes. The model does not address the evolu-

tion of these mechanisms themselves, but instead assumes that they are available (at a cost).

As we have already briefly mentioned, we should only expect to see a benefit to polyphenism if each alternative phenotype can find advantageous conditions within the population's environment or niche. To simulate this we vary the environmental conditions over time (although the incorporation of a spatial component with a corresponding temperature or chemical gradient would be one alternative). Thus, we expect to find that in a rapidly varying environment populations maintaining multiple alternative phenotypes would be at an advantage. Conversely, in a static environment we expect that this type of genotype will be selected against, due to the constraints and costs of unnecessarily supporting multiple phenotypes. We investigate both of these scenarios and find evidence to match our expectations supporting the alternative phenotypes hypothesis. Furthermore, the inverse relationship between polyphenism benefit and environmental stability indicates the potential for a more rigorous treatment of the APH.

In the next section, we provide further background on the APH. In Section III, we describe the evolutionary model that is used to support our claims. In Section IV, we describe the experiments performed, and provide results. A discussion is presented in Section V and finally future avenues for research are outlined.

The Alternative Phenotypes Hypothesis

West-Eberhard proposes a conceptual framework to aid the understanding of the role of phenotypic plasticity in providing organic innovation that can ultimately result in speciation. Here we outline the stages that comprise this framework of the alternative phenotypes hypothesis:

1. Prior to alternative phenotypes: the entire population/species exhibits a single phenotype
2. A switch mechanism arises in the population providing the capability for the context-sensitive expression of phenotypes
3. Novel alternative phenotypes evolve and stably persist in the population
4. Each alternative is subject to improvement by natural selection and becomes more specialist to some set of environmental conditions (the context-sensitive expression can prevent the alternative phenotypes from competing with one another)
5. Conditions in some locale may change to favour one alternative over the others; this phenotype will now become exclusively expressed
6. Character release: the genotype no longer has to support all of the alternative phenotypes

7. Accelerated speciation from parent population

A switch mechanism is required to determine which phenotype develops, and this could be allelic (genetic), condition-sensitive (environmental), or a combination of these factors. For the purposes of the APH it is not important which type of switch gives rise to the multiple phenotypes.

After a switch is established each phenotype can be evolved in semi-independence, but this need not lead to reproductive isolation. Since each of the phenotypes will only be selected on when they are expressed, each is buffered from negative selection provided that their expression is cued by the environment in which they are advantageous.

West-Eberhard's ideas on the significance of environmental influence on evolution are generally well accepted (see, e.g. Moran, 1992; van Buskirk, 2002; Bourke and Franks, 1991), although light criticism is given with respect to the sparsity of underlying mechanisms (Schlichting, 2003). Comprehensive evidence supports several of the stages contained within the alternative phenotypes hypothesis, including the existence of alternative phenotypes, phenotype fixation, character release and resulting speciation. However some aspects of the APH are not so clear, such as the identity of mechanisms that can provide alternative phenotypes, and the conditions that afford a selective advantage to populations that exhibit these mechanisms. The stages that have the largest evolutionary impact, speciation and macroevolution, can only proceed if conditions exist that are favourable for alternative phenotypes to stably exist. Thus, we limit the scope of this paper to address the investigation of stage 3: the conditions that favour polyphenic populations.

Examples of Alternative Phenotypes in Nature

There are many examples of species that exhibit alternative phenotypes, including the different castes of social insects (Wilson, 1971) with widely varying lifespans across these different castes (Jemielity et al., 2005), reproductive strategies of males (e.g. Spinney et al., 2006), and seasonal polyphenisms in aphids (Tauber et al., 1986). We focus on one single example as it provides an excellent illustration of several stages of the alternative phenotypes hypothesis: the buttercup *Ranunculus flammula* as studied by Cook and Johnson (1968). These plants can develop either lanceolate or linear leaves, when terrestrial or immersed in water, respectively. If plant populations are found in either permanently aquatic or permanently terrestrial conditions, they will only develop one type of leaf, but it is possible for a single plant to develop leaves of both types in response to changing conditions (e.g. in a lake with 'seasonally fluctuating water levels'). Cook and Johnson also described experiments where plants from monomorphic populations were put in environments unlike their own. When compared with heteromorphic populations, the monomorphic survival ability is reduced. This indicates that where phenotype fixation

has occurred (stage 5), specialisation to that environment is in progress towards character release (stage 6) and speciation (stage 7).

A Model of Alternative Phenotypes in a Varying Environment

In this section we describe the model system used to simulate the evolution by natural selection of a population with the capacity to support environmentally cued alternative phenotypes.

Maintaining multiple phenotypes is only beneficial if each one can become specialised within a different environment. Accordingly, we provide two environments that switch over time, according to a fixed rate, as we believe that these environments are sufficient to motivate the formation of alternative phenotypes. Additionally, the environmental niches should have a significant amount of overlap, since the two resulting phenotypes are still part of the same species and will share many ecological requirements. This overlap indicates that an individual that has specialised to one niche will not be completely unfit in another niche (although it will obviously be outcompeted by a specialist in this second niche).

As such, we find it suitable to extend and modify the framework used by Kashtan and Alon (2005; 2007). These authors are interested in understanding the mechanisms that might explain the modularity observed in biological systems. Following Lipson et al. (2002) they investigate how a varying environment could bring about the evolution of modularity by simulating the evolution of electronic logic circuits towards a target logical function F . They find that if the environment varies between modular functions, modular networks emerge despite an inherent cost: non-modular solutions that use fewer gates do exist. However, these were found to be much more vulnerable to the switching between environments. This is due to the target logic functions that form each environment: the functions share logical substructures and so good solutions to one target require only a small number of changes to satisfy the second target. The smaller, specialist, non-modular solutions would require significant re-working to satisfy the second target. As such, the more expensive modular solutions were favoured in modularly varying environments. We believe that the phenotypic distance between these two modular solutions can be reduced, by the APH, such that it can be reliably traversed within a single generation.

Adapting this work we define a target function F that switches between two states F_1 and F_2 , defined as follows:

$$F_1 = AND(XOR(I_1, I_2), XOR(I_3, I_4)) \quad (1)$$

$$F_2 = OR(XOR(I_1, I_2), XOR(I_3, I_4)) \quad (2)$$

with I_1 to I_4 representing Boolean inputs. Note that for 50% of the input patterns these two functions return the same value.

This could be thought of as, for example, a model of seasonal variation in which both photoperiod and temperature influence the development of the *Pieridae* family of butterflies (Shapiro, 1978). Alternatively, it could correspond to the environmental factors of diet, temperature and pheromones which determine the caste of social ants and other insects (Wheeler, 1986). Clearly, the specified target functions represent a substantial simplification of these examples. As the switch between these two functions is periodic we are abstracting away from the complex instability found in biological systems. However, since our representation for individuals has no capacity to learn this periodicity we feel that it is suitable for our purposes.

Given this environmental set-up, a population of individuals of size p , evolves using a generational genetic algorithm (Mitchell, 1996), towards a solution for each function in turn. Each individual G is represented by two sets of integers A and W . Formally these are constrained as follows:

$$A = \{x \in \{0, 1, 2\}\} \quad (3)$$

$$W = \{(y, z) | y, z \in \mathbb{Z}; -j \leq y, z \leq l\} \quad (4)$$

Each gene, g_i drawn from an individual consists of two linked parts x_i and $(y, z)_i$. Each value x represents the response of an individual gene to the environment. If $x = 0$ it will be expressed in either environment, if $x = 1$ it will be expressed in response to F_1 , and correspondingly if $x = 2$ it will be expressed in response to F_2 . When expressed each gene forms a NAND logic gate with index i , with two inputs connected to the gates denoted by y and z . The range of y and z is limited by the number of genes in the genome, l , and the number of inputs j . Accordingly, if y or $z < 0$ it will denote a connection to a corresponding input. Experiments in this paper use $j = 4$. Finally, if y or $z = 0$ it will denote a connection to the output of the output gate.

The inclusion of a set of environmental switches, A , and a variable length genome, l , represents a significant but necessary departure from Kashtan and Alon (2005). In limiting our changes to the introduction of alternative phenotypes we hope to couple the clear benefits of Kashtan and Alon's work to our own investigation of the APH.

To assess the quality of a given individual we determine its accuracy at solving a particular target logic function, F_1 or F_2 . This is performed by applying all 2^4 input patterns, from I_1 to I_4 , to the logic circuit defined by the expressed phenotype of an individual. The resulting output is then recorded and the proportion that matches the target function forms the initial fitness of the individual. To limit genetic growth and impose a penalty on the maintenance of alternative phenotypes this initial fitness value, f_{init} , is modified accordingly:

$$f_{final} = f_{init} \times 0.99^{(l-12)} \quad (5)$$

Having assigned a final fitness value to all members of the population, the next generation is constructed. It is composed of three parts: the fittest s individuals, S^+ , are copied directly without modification; the set of the least fit s individuals, S^- , is replaced by a copy of S^+ ; the remaining $(p - 2s)$ individuals inherit their genes from a randomly selected member of S^+ , subject to mutation at a rate of m . No crossover mechanism is used. Three events are possible if a mutation occurs:

1. **Gate change** — 50% of all mutations. A random gene, g_i is selected, and one of its inputs, y_i or z_i , is randomly changed with uniform probability within the range $-j \leq y, z \leq l_G$.
2. **Length change** — 30% of all mutations. With equal probability, a new random gate is added to the individuals genotype ($l_G + 1$) or a random gate is removed ($l_G - 1$).
3. **Environment sensitivity change** — 20% of all mutations. A random gene, g_i is selected, and its environmental switch, x_i , is changed with equal probability according to $x \in \{0, 1, 2\}$.

When initialised a population of size p is created, each individual within this population has a genome of length l_{init} and for all environmental switches $x = 0$. The simulation is then run for a fixed number of time steps t with the environment switching between the two logical functions F_1 and F_2 at a fixed rate r . In the following simulations each configuration is run 30 times, from which a number of averages are taken. Specifically, we record and plot the mean and maximum final fitness values, f_{final} , at each time step and the mean number of environment specific switches, $x = 1$ or $x = 2$, within each genome. If there is no selection pressure we can calculate the frequency with which environment specific switches will arise:

$$P(p - 2s) \times m \times P(x \in 1, 2) \times 0.2 \quad (6)$$

With 0.2 representing the fixed probability of an environment sensitivity change. Using the specified experimental parameters we would expect these switches to arise with a 5.2% frequency. We will now detail the specific parameter settings for our experimental work.

Simulated Experiments

Here we provide information regarding the experiments performed using the model described above, and detail their results.

The parameter settings that are common to all experiments described in this section are as follows: $m=0.5$, $p=100$, $l_{init}=13$, $s=10$, and $t=50,000$. All experiments were replicated 30 times.




Figure 1: A reproduction of Kashtan and Alon’s work with the addition of alternative phenotypes and a variable length genome. These modifications inflict a fitness penalty

Experiment 1: Introduction of Capacity for Alternative Phenotypes

Initially we test how the introduction of alternative phenotypes modifies Kashtan and Alon’s model, by reproducing their experiments with the environment switching at $r=20$.

Figure 1 shows the mean and maximum fitness of the polyphenic population, averaged across all 30 repeats. We note the mean fitness of the population moves between values of 0.55 and 0.75 throughout the experiment. This is somewhat contrasting to Kashtan and Alon’s system, where this population can evolve to an ideal solution for each environment between each switch. This provides us with some indication of how disruptive the inclusion of this plasticity mechanism is.

Experiment 2: Fixed Environment

An expectation stated in section I is that no selective advantage is conferred on a polyphenic population unless the environment varies. Thus we consider the case where environmental conditions are kept fixed throughout the experiment: r is set to t and the target function is set to F_1 . Figure 2 (a) shows the number of environment-sensitive loci in these static conditions. For comparison, frame (b) shows the same measure when $r=1$. We observe that in the static environment the population does not evolve a significant number of environment-sensitive genes. The maximum in the population is between 1 and 4, and the mean occasionally moves above zero due to drift. Alternative phenotypes are evidently not a strong feature of the population under static conditions. This behaviour contrasts with the results from the rapidly varying environment, where the mean number of environment-sensitive loci is around 2, and frequently the entire population has at least one such locus. This experiment was also duplicated with the target function held at F_2 ,

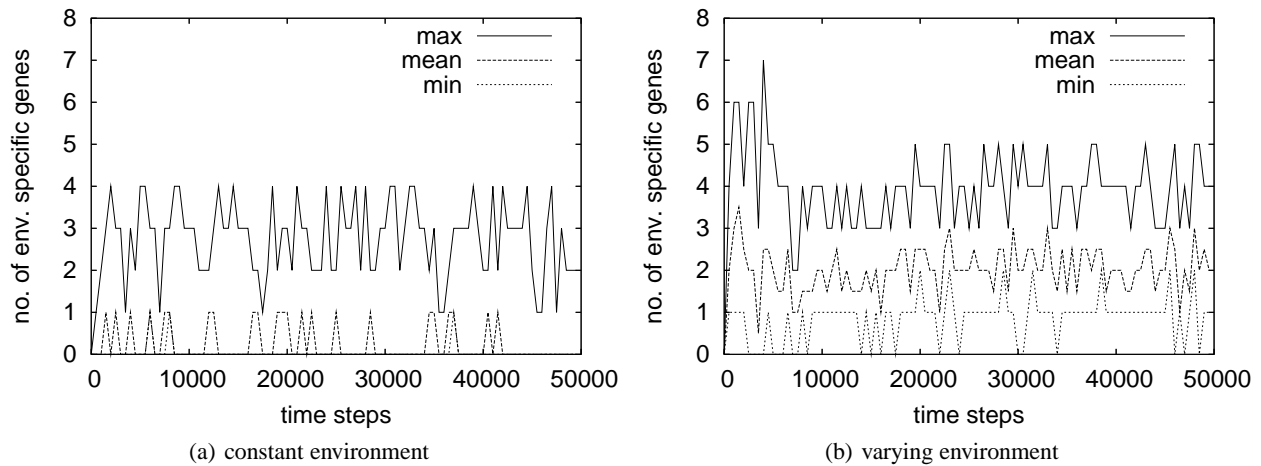


Figure 2: The effect of a rapidly switching environment compared to a static environment. In a static environment, frame (a), the population is almost entirely made up of specialists without any environmentally sensitive genes. By comparison, in a rapidly switching environment (frame (b)) alternative phenotypes are the norm

with qualitatively equivalent results.

Experiment 3: Rapidly Varying Environment

To investigate the potential benefits of maintaining alternative phenotypes in rapidly varying environments, we compare a population that has this capability with a population that does not. The environment is set to switch in each generation (*i.e.* $r=1$). In the monophenic population, all genes in an individual will be expressed regardless of the environment state, and mutations are restricted to modifying the phenotype length or gate assignments.

Only the population with polyphenic capability finds high fitness genotypes when in this rapidly varying environment. Figure 3 (a) shows the results from all 30 repeats of this experiment. We see that the population contains individuals of fitness of approximately 0.9, and the population mean is approximately 0.8 by the end of the experiments. This behaviour contrasts with the results shown in Figure 3 (b), which depicts the evolution of a population without polyphenism. The results indicate a population fluctuates around a fitness of 0.75 at best, and a significantly larger gap between the mean and maximum fitness values (the mean fitness fluctuates between 0.375 and 0.55). Consideration of $F1$ and $F2$ shows that for 50% of the input patterns these two functions return the same value, partially accounting for this mean fitness fluctuation.

To verify that environmentally cued loci are responsible, we consider Figure 2 (b). This illustrates the number of loci that will only appear in selected environments, for the polyphenic population in the rapidly varying environment. The mean number is around 2 and some individuals have many more (up to 5 in the long term). For the experiments in a static environment, we saw the results in Figure 2 (a).

Here the population is largely made up of individuals with no environmentally cued loci.

The results plotted in Figure 3 are for the most disruptive configuration that we tested, where the switching occurred at every generation. Experiments with $r=2$, $r=5$, and $r=10$ behave qualitatively similarly, although the highest fitness discovered degrades with increasing switching period. When we slow the environment switching to $r=20$, the behaviour moves towards that of the monophenic population. However, the polyphenic population mean fitness is much closer to the maximum fitness found than for the monophenic population.

Discussion

The results from experiments performed provide a perception on some of the conditions that affect the adaptive merit of maintaining alternative phenotypes. Experiment 1 reveals the cost of the plasticity that has been incorporated into the model. Experiment 3 demonstrates a clear advantage to populations maintaining alternative phenotypes when in an unstable environment, in spite of these additional costs. The polyphenic population has significantly higher fitness than the monophenic population in the most rapidly varying environments. Additionally, the mean and maximum fitness of the polyphenic population are a lot closer indicating a more stable region of adaptive space had been reached and the population had converged. When the environmental switching rate is reduced, the fitness improvement become less pronounced. However, when the fitness improvement is negligible at $r=20$, we still see the persistence of environment-sensitive genes (see Figure 4). There comes a point where maintaining alternative phenotypes is no longer viable as is shown in the limit case in experiment 2. This test confirmed

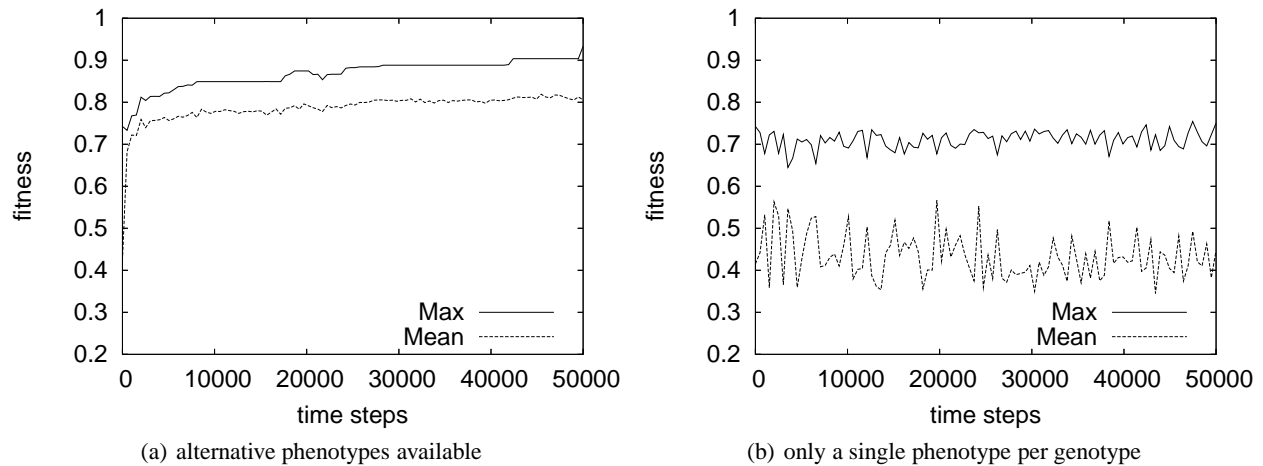


Figure 3: A population with the capacity to support alternative phenotypes compared against a population that cannot. The population in frame (a) is far more successful in handling the rapidly varying environment than the population in frame (b)

that polyphenism does not evolve in a fixed environment. Although we note that some drift is possible in the number of environment-sensitive genes: when stably in environment 1, an environment switch set to 1 will have the same effect on the phenotypic expression as that locus being set to 0 (and equivalently for environment 2). It is also worth considering the selection pressures on genes set to be expressed in environment 2 when it is fixed in environment 1. Although any such genes will not be expressed and thus cannot have a negative impact on the phenotype, there is nevertheless a cost to maintaining this gene. There is no pressure to maintain such a gene, and so we should expect these genes to be purged. Figure 2 does not show the values independently, but we can report that in experiment 2 all environmentally-cued genes match their respective environments (barring for a single generation in one anomalous result).

When considering the conditions in experiment 3, the population initially contains genotypes with low fitness to either environment. However, selection favours genotypes that express phenotypic traits that can contribute to high fitness in both environments, such that the population will move to a portion of the fitness landscape that overlaps. This can only be the case when the environments share a significant portion of their structure — and the target functions chosen by Kashtan and Alon have exactly this property.

There are sets of genes that co-occur due to the environmental cuing, and these sets are buffered from one another, forming an interesting parallel to evolutionary computing. A problem often faced in evolutionary algorithms, known as premature convergence, is when population diversity is lost rapidly. This can lead to the population converging on low fitness optima, and attempts to alleviate this are known as diversity maintenance (see Singh and Deb, 2006). These typically restrict the competition between individuals such

that portions of the population can focus on different parts of the fitness landscape. The condition-sensitive portions of a genome could potentially inspire a new diversity maintenance technique. Because one set of genes will only be expressed when in an environment that is advantageous for that particular phenotype, direct competition between alternatives is avoided.

A study into evolvability by Earl and Deem (2004) uses a ‘DNA swap’ mechanism that makes large, but non-random genetic changes in addition to small-scale changes by mutation. The DNA swap involves the substitution of genetic material for a particular genetic subdomain from a pool of low-energy alternatives for that subdomain. This can be considered as a form of diversity maintenance: the pools contain many different options to be swapped in and the alternative selected for the current environment, restricting competition being restricted to the subdomain (contrast this with the buffering provided by environmentally sensitive gene expression to a subset of genetic material in a single individual). They also investigate the suitability of each mechanism across a range of rates of environmental change, and find that large-scale variation is favoured increasingly in rapidly varying environments, further supporting the position that mutation alone is inadequate to cope with unstable conditions.

The exploration of a buffering-driven diversity maintenance mechanism for evolutionary computing is outside the scope of the current body of work, but considering condition sensitive switching in this light may help us to better understand the types of environment that it may prove advantageous within.

In our model all genes have the potential to be conditional on environmental cues, and this in principle allows several different configurations:




Figure 4: A population with a switching rate of 20 gains no fitness benefit from alternative phenotypes, see Figure 1. However, switching alone is sufficient for the persistence of alternative phenotypes within the population

- specialists: ideally adapted to a single environment; no condition-sensitive genes present
- modular generalists: suited to relatively slowly switching environment, where there is a region in the genetic space that has solutions for each environment nearby
- generalists with alternative phenotypes: suited to a faster-switching environment, where fit configurations for each niche overlap somewhat; consequently many genes are expressed in all environments and a small proportion of genes are environmentally sensitive
- polyphenic specialists: best suited to an environment that varies rapidly, but there is significant distance between fit phenotypes for each set of conditions; consequently an almost independent phenotype is supported for each environment

However we only encounter the first and third of these possibilities with the experimental conditions performed so far. One can imagine a line of inquiry that explores the conditions sufficient to give rise to the currently unseen genotype structures, and the potential trajectories between each of these types.

There is a heavy penalty associated with the number of genes required to code for two fit specialists, so it is unsurprising that they don't appear in the experiments performed to date. However a pair of target functions that varied in a non-modular way (*i.e.* with very little overlap) may give rise to such genotypes. This would still require two sets of genes to be co-adapted, and the isolation provided by the environment-sensitivity could maintain the correct diversity to lead to such adaptation.

Since selection in our model preserves the fittest individuals from each generation, one might expect the highest fitness to increase, or at least stay constant. We do observe fitness decreases in Figure 1, in contradiction with this expectation. This is because the high fitness solutions that are preserved from one environment to the next are not necessarily of high fitness in this new environment. They consequently may not be preserved long enough to return to high fitness in the second environment. However as shown in Figure 3, the inclusion of environmentally-switched gene expression goes a long way towards mitigating this problem.

We will now outline a number of aspects for future research. Initially, we would like to understand the disruption inflicted upon polyphenic populations, and the factors to which it can be attributed. We also wish to perform similar experiments to those reported here, but with different sets of target functions. As discussed above, this includes considering environments that do not have modular overlap with the aim of evolving more independence in the alternative phenotypes of one individual.

There are also avenues that involve extensions to the current framework. This includes connecting together additional stages of the APH. For example, investigating how long character release might take would be possible by changing the schedule of environmental conditions experienced by a population with established alternative phenotypes. Additionally, it would be valuable to test the hypothesis that an environmentally cued alternative phenotype can lie dormant for many generations without being expressed. The identification of candidate switch mechanisms would require significant extension to the model. We could investigate what conditions might enable useful cuing mechanisms to arise, using the framework that we have established to test the plausibility of a particular switch mechanism. Finally, we could employ a slightly different experimental set up to study direct competition between polyphenic and monophenic populations instead of comparing the populations in isolation.

As far as we know, we have presented the first individual-based simulation model of a portion of the alternative phenotypes hypothesis, illustrating some capabilities and limitations of polyphenism under different abstract environmental conditions. The results that we have obtained so far, whilst modest, provide support for the later stages of the APH that require the stable existence of alternative phenotypes to proceed. We feel that there are many aspects that could be better understood when using simulation models to enrich thought experiments (Di Paolo et al., 2000). In developing this initial model, we have identified abstract environmental conditions that could account for the previously observed phenomena within the APH. We believe that the conceptualisation provided by this approach could ultimately unify the current experimental evidence into an increasingly rigorous underlying framework.

References

- Barandiaran, X. and Moreno, A. (2006). Alife models as epistemic artefacts. In Rocha, L. M., editor, *Procs ALifeX*, pages 513–519.
- Bourke, A. F. and Franks, N. R. (1991). Alternative adaptations, sympatric speciation and the evolution of parasitic, inquiline ants. *Biological Journal of the Linnean Society*, 43(3):157–178.
- Cook, S. A. and Johnson, M. P. (1968). Adaptation to heterogeneous environments. I. variation in heterophylly in *Ranunculus flammula* L. *Evolution*, 22(3):496–516.
- Dennett, D. C. (1994). Artificial life as philosophy. *Artificial Life*, 1(3):291–292.
- Di Paolo, E., Noble, J., and Bullock, S. (2000). Simulation models as opaque thought experiments. In *procs Artificial Life VII*, pages 497–506, Cambridge, MA. MIT Press.
- Earl, D. J. and Deem, M. W. (2004). Evolvability is a selectable trait. *PNAS*, 101(32):11531–11536.
- Geard, N. L. and Wiles, J. (2005). A gene network model for developing cell lineages. *Artificial Life*, 11(3):249–268.
- Hall, B. K. (1998). *Evolutionary developmental biology*. Chapman & Hall, London, 2nd edition.
- Hornby, G. S. and Pollack, J. B. (2001). Body-brain co-evolution using L-systems as a generative encoding. In Goodman, E. D., editor, *Procs GECCO*, pages 868–875. Morgan Kaufmann.
- Jemielity, S., Chapuisat, M., Parker, J. D., and Keller, L. (2005). Long live the queen: studying aging in social insects. *AGE*, 27(3):241–248.
- Kashtan, N. and Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *PNAS*, 102:13773–13778.
- Kashtan, N., Noor, E., and Alon, U. (2007). Varying environments can speed up evolution. *PNAS*, 104(34):13711–13716.
- Lande, R. (1980). Sexual dimorphism, sexual selection, and adaptation in polygenic characters. *Evolution*, 34(2):292–305.
- Lipson, H., Pollack, J. B., and Suh, N. P. (2002). On the origin of modular variation. *Evolution*, 56(8):1549–1556.
- Mills, R. and Watson, R. A. (2006). On crossing fitness valleys with the Baldwin effect. In Rocha, L. M., editor, *Procs ALifeX*, pages 493–499.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press.
- Moran, N. A. (1992). The evolutionary maintenance of alternative phenotypes. *American Naturalist*, 139(5):971–989.
- Nolfi, S. and Parisi, D. (1998). Neural networks in an artificial life perspective. In *Procs ICANN*, pages 733–737.
- Schlichting, C. D. (2003). Environment rules. *Trends in Ecology and Evolution*, 18(10):496–497.
- Shapiro, A. M. (1978). The evolutionary significance of redundancy and variability in phenotypic-induction mechanisms of pierid butterflies (Lepidoptera). *Psyche*, 85(2–3):275–283.
- Singh, G. and Deb, K. (2006). Comparison of multi-modal optimization algorithms based on evolutionary algorithms. In *Procs GECCO*, pages 1305–1312.
- Spinney, L. H., Bentley, G. E., and Hau, M. (2006). Endocrine correlates of alternative phenotypes in the white-throated sparrow (*Zonotrichia albicollis*). *Hormones and Behavior*, 50(5):762–771.
- Tauber, M. J., Tauber, C. A., and Masaki, S. (1986). *Seasonal Adaptations of Insects*. Oxford University Press.
- van Buskirk, J. (2002). A comparative test of the adaptive plasticity hypothesis: Relationships between habitat and phenotype in anuran larvae. *American Naturalist*, 160(1):87–102.
- West-Eberhard, M. (1986). Alternative adaptations, speciation, and phylogeny (a review). *PNAS*, 83(5):1388–1392.
- West-Eberhard, M. (1989). Phenotypic plasticity and the origins of diversity. *Annual Review of Ecological Systems*, 20:249–278.
- West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. Oxford University Press.
- Wheeler, D. E. (1986). Developmental and physiological determinants of caste in social Hymenoptera: Evolutionary implications. *American Naturalist*, 128(1):13–34.
- Wilson, E. O. (1971). *The Insect Societies*. Harvard University Press, Cambridge, MA.
- Wolpert, L. (2007). *Principles of Development*. Oxford University Press, 3rd edition.