

Committees as Artificial Organisms – Evolution and Adaptation

Roberto C. Alamino

Aston University - Mathematics
Birmingham - UK

Abstract. Generalised committee machines are here proposed to model the interaction of the DNA of a simplified artificial organism with its environment and shown to induce a unique genotype-phenotype map. An application to organisms being subjected to a toxic environment is shown to allow for a generalised form of antagonistic pleiotropy. The same scenario is studied in order to show the difference in adaptation in the presence of a fitness cost given by a lower reproduction rate.

1 Introduction

Natural selection was life's first strategy on Earth to transmit information about the environment to following generations, mostly through genes. Learning happens mainly by random mutations, although it also involves other processes like sexual reproduction, chromosomal crossover and horizontal gene transfer.

The interaction between genes and the organism's environment is not direct, but mediated by cascades of physico-chemical processes. Although mutations happen on the gene level (*genotype*), it is the macroscopic characteristics generated by them (*phenotype*) that suffer the selection pressure. The map between these levels is usually not 1-1 and is called a genotype-phenotype map (GPM).

Here, a model is introduced in which the interaction with the environment is mediated by committee machines whose units are perceptrons. These *artificial organisms* live on an idealised lattice representing their habitat. This model, introduced in section 2, allows for adaptation via natural selection and can be tuned to account for several biological mechanisms. Section 3 shows that there is a *unique* GPM induced by it and provides its main properties. In section 4 we introduce a simplified scenario of an organism in a toxic environment and show that it can present a generalised form of antagonistic pleiotropy. In section 5 the influence of the fitness cost of a smaller reproduction rate is analysed.

2 Artificial Organisms

Consider a square $N \times N$ lattice with periodic boundary conditions in both directions in which the artificial organisms live. A random initial configuration of occupied and empty sites will be evolved in time with the objective of following the dynamics of a proxy for the quality of the organisms' adaptation.

The organisms cannot move, but can reproduce. Here, only the case of asexual reproduction is considered – at each time step t , each organism has a probability r of cloning itself to one of its *empty* neighbouring sites with equal

probability. For simplicity, organisms possess a single DNA strand, represented by a binary chain $\pi_{ij} \in \{\pm 1\}^D$, where D is an integer, with a probability m per base of mutating upon reproduction.

The response to the environment is given by a *death probability* q_{ij} , the probability that the organism at the i -th row and j -th column dies due to the local environmental conditions. A minimal requirement for a GMP is that it has to allow for *learning* and *forgetting*, both essential for adaptation. Forgetting is important in dynamical environments as it allows acquired adaptations to fade away if they become non-competitive or harmful. One of the most basic models with these characteristics is the *perceptron* [1, 2].

The perceptron is the simplest machine learning model. It is characterised by a function of a multidimensional vector into a number, the *activation function*, which is usually a function of the scalar product of its *synaptic vector* (a vector encoding learnt information) and the *input vector* (encoding external stimuli).

The organisms change their output according to external stimuli, encoded by binary chains $\mathbf{A} = (\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^n)$, with $\mathbf{a}^\mu \in \{\pm 1\}^D$ and $\mu = 1, \dots, n$, where n is the number of perceptrons in the committee representing the organism.

Consider n perceptrons with activation functions $\phi_\mu(\langle \mathbf{a}^\mu, \boldsymbol{\pi} \rangle)$, $\mu = 1, \dots, n$, where $\langle \mathbf{x}, \mathbf{y} \rangle \equiv \mathbf{x} \cdot \mathbf{y} / D$. The response of the organism to the environment is

$$q = \Phi(\phi_1(\langle \mathbf{a}^1, \boldsymbol{\pi} \rangle), \dots, \phi_n(\langle \mathbf{a}^n, \boldsymbol{\pi} \rangle)),$$

where lattice indices were dropped for convenience. While no restriction is imposed on the range of the activation functions, Φ – the integration function of the committee machine – is required to be real-valued on $[0, 1]$ as it is a probability. We will often drop the arguments of the activation functions for convenience.

Although simple perceptrons cannot approximate general functions, adding one extra *layer* turns them into universal approximators [3]. For this work, one layer suffices, but the model is flexible and allows for more complex structures.

The form of Φ and dimension of \mathbf{A} can now be justified. They were chosen to allow mutations in the genotype to simultaneously affect all interactions. This corresponds roughly to the phenomenon known in genetics as *pleiotropy* [4]. In nature, each protein often participates in several metabolic process simultaneously and a single gene mutation can affect more than one of them. The difference here is that we are looking at a kind of pleiotropy at the base level, which is not the usual meaning in biology. However, if the coordinates of $\boldsymbol{\pi}$ are considered to represent the genes themselves (for instance, as in [5]), the two definitions become the same.

The normalisation of the scalar product also requires justification as it leads to a unique GPM. The choice allows for general activation functions to be $O(1)$ in D , which is the case in biological realistic scenarios.

3 The Genotype-Phenotype Map

In the simulations, the environment is the same everywhere and generated randomly with the components of \mathbf{a}^μ being i.i.d. variables. Straightforward but

lengthy calculations show that, for large D , the average values of the corresponding activation functions are approximately $\langle \phi_\mu \rangle \approx \phi_\mu(\bar{a}^\mu \bar{\pi})$, where \bar{a}^μ is the mean of the components of a^μ and $\bar{\pi} = D^{-1} \sum_{i=1}^D \pi_i$.

Although environments can differ microscopically, the committee can then only sense them as points in an n -dimensional space. A crucial point in evolutionary biology is to find appropriate GPMs [6, 7]. Here, a natural GPM is induced by the genome interaction with the environment and is given by $\Pi(\pi) = D\bar{\pi}$.

This map is an *emergent property*, a collective effect of a large number of bases. From the structural properties of GPMs considered in [7], the map Π possesses *redundancy*, as it is a many-to-one map which maps the set $\{\pm 1\}^D$ with cardinality 2^D onto the set $\{-D, -D+1, \dots, D-1, D\}$ with cardinality $D+1$ (an exponential reduction in size) and *bias*, meaning that the number of genotypes for each phenotype is not the same: given a certain dimension D and phenotype Π , the number of genotypes corresponding to it is simply the binomial coefficient

$$\binom{D}{(D+\Pi)/2}.$$

There are two other properties generally required, *robustness* and *evolvability*. Simulations show that the map allows for adaptation, but it is not robust. The pre-images in genotype space of the phenotypes are not connected, meaning that single mutations necessarily change the phenotype. On the other hand, the network of phenotypes is fully connected by single mutations, allowing any phenotype or genotype to be reached from any other. Although the map is not robust in principle, for large D , neighbouring phenotypes result in values of the activation functions that differ typically by small quantities, which can be seen as a *quasi-robustness*. Single mutations then lead to small phenotype variations and, unless the fitness landscape becomes discontinuous in a significant number of points, this should not affect evolvability.

4 Antagonistic Pleiotropy

The general model allows for several features from real organisms to be analysed. The simple case introduced below models the effect of a toxic substance (from here on, *poison*) present in the environment. Consider $n = 2$ and the same functional form for the death probability of all organisms

$$q = \Theta(\Delta)(1 - e^{-\phi_2 \Delta}), \quad \Delta = c - \phi_1,$$

where $\Theta(\Delta) = 0$ if $\Delta < 0$ and 1 otherwise, i.e., the poison is only harmful if its concentration c is above a certain threshold dependent on the organism's genotype. The function ϕ_2 is a sensibility – how much the death rate increases with the poison's concentration.

One possible form for the activation functions is given below

$$\phi_1(y) = \phi_2(y) = \frac{1+y}{1-y},$$

and $a_i^1 = a_i^2 = -1$, $i = 1, \dots, D$, with $\phi_1, \phi_2 \in [0, \infty)$. One can rewrite it as $\phi_1 = \phi_2 = x/(1-x)$, with $x = n/D$ and n the number of DNA coordinates equal to -1. When $D \rightarrow \infty$, x becomes a continuous variable on $[0, 1]$.

Although both activation functions are the same, the two perceptrons play antagonistic roles. While a higher ϕ_1 improves the organism's resistance to the poison, a higher ϕ_2 decreases it by making the organism more sensitive to it. This phenomenon in which the same genes encode one beneficial and one detrimental phenotype is known as *antagonistic pleiotropy* [8], where once again we are using the broader interpretation of the term discussed previously. It is important to notice that this is a property induced by the environment and not a feature of the organism's genotype only. The same activation functions in a different environment might not be antagonistic as, for instance, if $a_i^2 = 1$.

5 Fitness Cost

Beneficial mutations can have fitness costs – disadvantages with respect to the *wild-type* (the non-mutated genotype) – which might be a problem if the environment changes. There is evidence that this is the case for antibiotic resistance, where the fitness cost becomes a reduced growth rate [9]. To study this case, let us define the distance between a genotype $\boldsymbol{\pi}$ and the wild-type by their *Hamming distance* – the number of mismatching coordinates in the corresponding vectors. The wild-type generally corresponds to the typical genotype of the most adapted organisms. Here we arbitrarily use the genotype $\mathbf{1}$ with all coordinates equal to 1, which result in $H(\boldsymbol{\pi}, \mathbf{1}) = \frac{1}{2}(D - \Pi)$.

A precise experimental characterisation of the fitness cost is difficult, but a reasonable qualitative behaviour can be obtained by an exponential decay with a Boltzmann factor given by $r = r_0 e^{-\omega H(\boldsymbol{\pi}, \mathbf{1})}$, where r_0 is the reproduction rate of the wild-type and $\omega > 0$ is a positive decay constant controlling the decrease of the reproduction rate with the distance to the wild-type.

The lattice is initialised by putting an organism in each site with probability 1/2. Genotype bases are drawn independently with probability 1/2. The dynamics follow two steps at each t : (1) *Reproduction+Mutation* – all organisms are drawn once in a random order; each has a probability r of spreading a mutated clone to one of its empty neighbours with a probability m of *mutation* per base (a flip). (2) *Death* – all organisms are checked against q_{ij} for dying.

Each run of the simulation consists of T time steps where the double average of the death probability q_{ij} – over all organisms and initial configurations – and the population are recorded. To avoid stalling due to lack of space, 1/2 of the organisms are wiped out at random if the network becomes full.

The poison concentration is chosen such that the initial q_{ij} is 0.3 for the phenotype $\Pi = 0$. This value, constant during the simulation, allows time for adaptation before the population is wiped out. The environment is also fixed and chosen to allow adaptation, which is not always the case. Because $D = 100$ and $L = 50$, one can use the approximation for large D and fix the environment by the average values $\bar{a}^1 = 0.4$ and $\bar{a}^2 = 1$. The phenotype of the initial population




Fig. 1: The plots show the average death probability q for (A) $\omega = 0$ and (C) $\omega = 0.01$ together with the size of the populations for the same values of ω , respectively at (B) and (D), at time $T = 5001$ for different values of the mutation rate m .

is binomially distributed, broad enough to result in adaptation even if $m = 0$.

Because $q_{ij} \in [0, 1]$, instead of using the standard deviation for error bars, the deviations below and above the mean are separately calculated. Let \bar{q}_k be the average death probability for the k -th run of the simulation and \hat{q} its average over all runs. The deviations above and below the mean for q are respectively

$$\sigma_{\pm}^2 = \frac{1}{n_{\pm}} \sum_k \left[\frac{1 \pm \Theta(\Delta q_k)}{2} \right] (\Delta q_k)^2,$$

where $\Delta q_k = \hat{q} - \bar{q}_k$ and n_+ (n_-) is the number of cases in which the variation is non-negative (negative). In practice, the difference to the variance is not significant except when q is close to the borders of its interval. An analogous definition is used for the deviations of the population.

Fig. 1 shows the value of the population and the death probability for the different values of m at $T = 5001$ for $\omega = 0$ (no fitness cost) and $\omega = 0.01$ averaged over 100 runs. Although the fitness cost causes the collapse of the populations above a certain value of m , which is not unexpected as they grow slower, this seems to be little effect on q (apart from a higher variability when the population disappears), which means that the quality of the adaptation does not change. A more detailed analysis is under way to clarify this behaviour.

6 Conclusions

The main contribution of this work is a new model for representing the interaction between the genetic code of an organism and its environment. The model induces a unique GPM in the limit of a large number of genetic bases. A simple application to an organism in a poisonous environment was shown to allow for antagonistic pleiotropy and was used to study the effect of a fitness cost to adaptation.

An interesting possibility is to use this model to find encodings for environmental conditions. Given that several DNA sequences encoding certain phenotypes are already known, one can use them to learn the environmental vectors. Once this is done, they can be used to make predictions of the interaction of different DNAs with an environment without actually introducing the organism.

Another interesting direction of research is the possibility of learning the integration function of the committee machine. Given it is a general real valued function, this requires techniques which are still not available and will be the subject of future research.

I would like to thank Dr J. Neirotti, Dr M. Stich, Dr M. Chli and Dr A. Cheong for their suggestions and discussions.

References

- [1] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386, 1958.
- [2] A. Engel and C. van den Broeck. *Statistical Mechanics of Learning*. Cambridge University Press, March 2001.
- [3] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.
- [4] Frank W Stearns. One hundred years of pleiotropy: a retrospective. *Genetics*, 186(3):767–773, 2010.
- [5] Christopher F Taylor and Paul G Higgs. A population genetics model for multiple quantitative traits exhibiting pleiotropy and epistasis. *Journal of Theoretical Biology*, 203(4):419–437, 2000.
- [6] Peter F Stadler and Bärbel MR Stadler. Genotype-phenotype maps. *Biological Theory*, 1(3):268–279, 2006.
- [7] Sebastian Edmund Ahnert. Structural properties of genotype-phenotype maps. *Journal of The Royal Society Interface*, 14(132):20170275, 2017.
- [8] George C Williams. Pleiotropy, natural selection, and the evolution of senescence. *Evolution*, 11(4):398–411, 1957.
- [9] Dan I Andersson and Diarmaid Hughes. Antibiotic resistance and its cost: is it possible to reverse resistance? *Nature Reviews Microbiology*, 8(4):260–271, 2010.