

# Tierra's missing neutrality: case solved.

Russell K. Standish

School of Mathematics, University of New South Wales, Sydney, 2052, Australia

R.Standish@unsw.edu.au, <http://parallel.hpc.unsw.edu.au/rks>

## Abstract

The concept of neutral evolutionary networks being a significant factor in evolutionary dynamics was first proposed by Huynen *et al.* about 7 years ago. In one sense, the principle is easy to state — because most mutations to an organism are deleterious, one would expect that neutral mutations that don't affect the phenotype will have disproportionately greater representation amongst successor organisms than one would expect if each mutation was equally likely.

So it was with great surprise that I noted neutral mutations being very rare in a visualisation of phylogenetic trees generated in *Tierra*, since I already knew that there was a significant amount of neutrality in the *Tierra* genotype-phenotype map.

It turns out that competition for resources between host and parasite inhibits neutral evolution.

Keywords: *Tierra*, neutral evolution, genotype-phenotype map, Vienna RNA package

## Introduction

The influence of *neutral networks* in evolutionary processes was first elucidated by Peter Schuster's group in Vienna in 1996 (Huynen *et al.*, 1996; Reidys *et al.*, 1997). Put simply, two *genotypes* are considered *neutrally equivalent* if they map to the same *phenotype*. A *neutral network* is a set of genotypes connected by this neutrality relationship on links with Hamming distance 1 (i.e. each link of the network corresponds to a mutation at a single site of the genome). It should be noted that this definition is subtly different from that employed in Kimura's *neutral evolution theory* (Kimura, 1983), as in that theory, neutrality is defined as equivalence of fitness values, a notion that is ill-defined in coevolutionary systems. However as phenotypically equivalent organisms are neutral in Kimura's sense when a fitness function exists, much of neutral theory can be carried over into discussion of phenotypic neutrality.

Schuster's group noted that evolution tended to proceed by diffusion along these neutral networks, punctuated occasionally by rapid changes to phenotypes as an adaptive feature is discovered. The similarity of these dynamics with the theory of Punctuated Equilibria (Eldridge, 1985) was noted

by Barnett (1998). It was also noted that if a *giant network* existed that came within a hop or two of every possible genotype, evolution will be particularly efficient at discovering solutions, since only a few non-neutral mutations are needed to reach the optimum solution.

Most work on neutrality in evolution uses the genotype-phenotype mapping defined by folding of RNA (Schuster *et al.*, 1994). This mapping is implemented in the open source Vienna RNA package<sup>1</sup>, so is a convenient and well-known testbed for ideas of neutrality in evolution.

Also in 1996, I developed a definition of the genotype-phenotype mapping for *Tierra*, which was first published in 1997 (Standish, 1997). I noticed the strong presence of neutrality in this mapping at that time, which was later exploited to develop a measure of complexity of the *Tierran* organism (Standish, 1999; Standish, 2003). In 2002, I started a programme to visualise *Tierra*'s phylogenetic trees and neutral networks (Standish and Galloway, 2002) in order to “discover the unexpected”. Two key findings came out of this: the first being that *Tierra*'s *genebanker*<sup>2</sup> data did not provide clean phylogenetic trees, but had loops, and consisted of many discontinuous pieces. This later turned out to be due to *Tierra*'s habit of reusing genotype labels if those genotypes were not saved in the *genebanker* database. This might happen if the population count of that genotype failed to cross a threshold. This is all very well, except that a reference to that genotype exists in the parent field of successor genotypes. The second big surprise was the paucity of neutral mutations in the phylogenetic tree. We expect most mutations to an organism to be deleterious, and so expect that neutral mutations will have disproportionately greater representation amongst successor genotypes than one would expect if each mutation was equally likely.

## Neutrality in *Tierra*

*Tierra* (Ray, 1991) is a well known artificial life system in which small self-replicating computer programs are ex-

<sup>1</sup><http://www.tbi.univie.ac.at/~ivo/RNA>

<sup>2</sup>The *genebanker* is a database in which *Tierra* stores the genotypes that arise during evolution.

*Copyrighted Material*

ecuted in a specially constructed simulator. These computer programs (called digital organisms, or sometimes “critters”) undergo mutation, and radically novel behaviour is discovered, such as *parasitism* and *hyperparasitism*.

It is clear what the genotype is in Tierra, it is just the listing of the program code of the organism. The phenotype is a more diffuse thing, however. It is the resultant effect of running the computer program, in all possible environments. Christoph Adami defined this notion of phenotype for a similar artificial life system called *Avida* (Adami, 1998). In *Avida*, things are particularly simple, in that organisms either reproduce themselves at a fixed replication rate, or don’t as the case may be, and optionally perform range of arithmetic operations on special registers (defined by the experimenter).

In Tierra, organisms do interact with each other via a template matching mechanism. For example, with a branching instruction like `jmpo`, if there is a sequence of `nop0` and `nop1` instructions (which are no-operations) following the branch, this sequence of 1s and 0s is used as a template for determining where to branch to. In this case the CPU will search outwards through memory for a complementary sequence of `nop0`s and `nop1`s. If the nearest complementary sequence happens to lie in the code of a different organism, the organisms interact.

To precisely determine the phenotype of a Tierran organism, one would need to execute the soup containing the organism and all possible combinations of other genotypes. Whilst this is a finite task, it is clearly astronomically difficult. One means of approximation is to consider just interaction of pairs of genotypes (called a tournament). Most Tierran organisms interact pairwise — very few triple or higher order interactions exist. Similarly, rather than running tournaments with all possible genotypes, we can approximate matters by using the genotypes stored in a genebanker database after a Tierra run. In practice, it turns out that various measures, such as the number of neutral neighbours, or the total complexity of an organism are fairly robust with respect to the exact set of organism used for the tournaments.

So the procedure is to pit pairwise all organisms in the genebanker against themselves, and record the outcome in a table (there is a small number of possible outcomes, which is detailed in (Standish, 1997)). A row of this table is a phenotypic signature for the genotype labeling that row. We can then eliminate those genotypes with identical signatures in favour of one canonical genotype. This list of unique phenotypes can be used to define pragmatically a test for neutrality of two different genotypes, that may have generated by mutation from genotypes recorded in the genebanker. Pit each organism against the list of unique phenotypes, and if the signatures match, we have neutrality. The soup is then

this experiment is available from the author’s website.<sup>3</sup>

Tierra has three different modes of mutation:

**Cosmic Ray** A site of the soup is randomly chosen and mutated;

**Copy** Data is mutated during the copy operation;

**Flaw** Instructions occasionally produce erroneous results

Furthermore, in the case of cosmic ray and copy mutations, a certain proportion of mutations involve bit flips, rather than opcodes being substituted uniformly. This proportion is set as a parameter in the `soup_in` file (`MutBitProp`) — in these experiments, this parameter is set to zero.

In order to study the issue of whether neutrality is greater or less than expected in Tierra, I generated three datasets with each of the 3 modes of mutation operating in isolation. The sizes of each data set was 69,139, 87,003 and 198,982 genotypes respectively, generated over a time period of about 1000 million executed instructions. Genebanker’s threshold was set to zero, so all genotypes were captured. This led to a proper phylogenetic tree. After performing a neutrality analysis, a set of 83, 86 and 158 unique phenotypes was extracted as the test set for the tournaments.

Since the neighbourhood size increases exponentially with neighbourhood diameter, I restrict analysis to single site, or point mutations. In each data set, around 7% of these genotypes were created by a mutation at a single site and were neutrally equivalent to its parent. For each of these, I compute the number of neutral neighbours  $n_i$  existing in the 1 hop neighbourhood of the parent genotype  $i$ . The 1 hop neighbourhood size is  $32^{\ell_i}$ , where  $\ell_i$  is the length of the genome. For a given parent  $i$ , the ratio

$$r_i = \frac{v_i 32^{\ell_i}}{o_i n_i} \quad (1)$$

gives the proportion of neutral links actually followed relative to the number of neutral links available (*neutrality excess*), where  $v_i$  is the number of neutrally equivalent offspring, and  $o_i$  the total number of offspring and  $n_i$  the size of the 1 hop neutral neighbourhood. Fig. 1 shows the running average of this quantity over these transitions, with the genotypes numbered in size order.

Since all daughter genotypes are recorded, no selection is operating. In this case, one would expect that the proportion of neutral variants seen should be identical to the proportion of neutral variants within the 1 hop neighbourhood, and hence the neutrality excess should be identical to 1. However, in the case of instruction and cosmic ray flaws, not every daughter genome will make it into the genebanker. In the case of instruction flaws, it is rather unpredictable what the effect is. In the case of cosmic ray mutations, 50% of

<sup>3</sup><http://parallel.hpc.unsw.edu.au/getaegisdist.cgi/getsource/ecology/3.D3>




Figure 1: Running average of neutrality excess ( $\langle r_i \rangle$ ). Genomes are ordered according to size, and neutrality excess is averaged over all genomes to the left of that data point. Three different datasets are analysed, with each of the three modes of mutation turned on. Then the datasets are further filtered to only include offspring whose maximum population count is greater than 1, i.e. selection is operating.

time one would expect the parent to be mutated, rather than the daughter. In the case of a mutation affecting a crucial gene of a parent genotype, the organism may not be able to reproduce at all, thus favouring neutral mutations. Only copy mutations should affect all sites of the genome equally, leading to a neutrality excess equal to one. The measured value, however is about 1.3, substantially greater than one. The reason for this is not known at this point in time.

The datasets were further subsetted to include just those transitions whose daughter organism successfully reproduced, i.e. with a maximum population count greater than 1. The neutrality excess in this case is substantially less than 1, so something in Tierran evolution is favouring non-neutral evolution.

### Competition Effects

Consider a single species ecosystem with logistic dynamics:

$$\dot{x} = rx(1 - x/K), \quad (2)$$

where  $x$  is the population size,  $r$  the net reproductive rate and  $K$  the carrying capacity. A phenotypically equivalent genotype attempting to invade this ecosystem will have the following dynamics:

$$\dot{x}' = rx'(1 - x'/K) \approx 0, \quad (3)$$

( $x'$  being the population size of the invading genotype) as  $x \approx K$  at equilibrium. So there is a substantial likelihood that the neutral variant fails to invade the ecosystem.

This argument is of course an extreme case. Stochastic effects due to finite population sizes will increase the chances of a neutral variant invading the ecosystem, however the point still remains that the neutral variant is not on an equal footing as the incumbent.

In Tierra, however, there is an age structure in the population, with organisms being placed in a *reaper queue*, from which the oldest organisms are selected when death is required. This fact alone implies that neutral variants of self-replicating organisms will successfully replicate, and hence cannot be responsible for the neutrality deficiency.

However, consider a Tierran ecosystem consisting of hosts and parasites, where the parasite require the presence of a host organism within a certain distance of itself in the soup, in order for the parasite to replicate. Since parasitic organisms replicate faster than the hosts (due to their smaller program lengths), they tend to displace host organisms until there are not enough hosts to go around. At which point, the parasite's fecundity drops. At equilibrium, the effective reproductive rates of host and parasite are equal.

A neutral variant will therefore be quite likely to not have a suitable host in its neighbourhood to allow it to replicate. Consequently, neutral evolution is suppressed amongst parasites. In the next section I will test this idea by setting up an artificial host-parasite coevolutionary system, using the well known RNA genotype-phenotype map.

### Vienna RNA Folding Experiments

It is quite well known that evolution using the RNA folding map (Schuster et al., 1994) exhibits a great deal of neutrality, at least for a standard genetic algorithm optimising a well defined fitness function. Until now, evolutionary systems based on the RNA map exhibit the unsurprising result of neutrality excess defined by eq (1) being greater than or equal to 1. I now present results of an RNA map experiment that demonstrates neutrality suppression ( $r_i < 1$ ), based on the resource competition explanation elaborated earlier. We need two types of organism (host and parasite) competing for a fixed space that can support  $N = 100$  organisms. Parasites can only reproduce if they are situated next to a host (neighbourhood size  $v = 2$ ), but reproduce twice as fast as the host type.

Once an organism has reproduced, it replaces the least fit organism. Fitness is determined by how close the parasitic phenotype is to any hosts in the neighbourhood of the parasites, and decreases in a similar way with the similarity of the parasites in the neighbourhood for host organisms:

$$F_h = 1 - \frac{1}{v\ell} \sum_{i \in \mathcal{P}_v} d(i, h) \quad (4)$$

$$F_p = \frac{\rho}{v\ell} \sum_{i \in \mathcal{H}_v} d(i, p) \quad (5)$$

where  $h$  and  $p$  are host and parasite genotypes respectively,  $\mathcal{H}_v$  and  $\mathcal{P}_v$  the set of hosts and parasites respectively within the neighbourhood of size  $v$  of  $p$  and  $h$  respectively.  $d(i, j)$  is the string edit distance between the phenotypes<sup>4</sup>, and  $\ell$

<sup>4</sup>The string edit distance is related to the Hamming distance




Figure 2: Neutrality excesses for RNA folding host-parasite system as described in the text.

is the gene length (set equal to 20 for all organisms in this experiment).

The factor  $\rho$  adjusts the relative dominance of parasites over hosts. Set it too low, and hosts will eliminate the parasites by virtue of replacing them when replicating. Set it too high, and hosts will only be competing with themselves. In this experiment, a value of  $\rho = 3.1$  was found to give intermediate behaviour.

An alternative version of this experiment where organisms were selected at random for death, rather than according to a fitness relationship showed similar dynamics, although the neighbourhood size  $v$  needed to be increased to 4 to allow a stable population of parasites to persist.

Figure 2 shows the neutrality excess for this experiment. The model will consistently produce a neutrality deficiency for the parasites over a broad range of model parameters. If the parameter  $\rho$  is set too high, the hosts will compete strongly with themselves, suppressing neutrality in the host population also.

Source code for this experiment is available from the author's website.<sup>5</sup>

## Conclusion

The suppression of neutrality in Tierran evolution is a real effect. An explanation couched in terms of host parasite competition was found, and a model was constructed using

(no. of base pairs that differ between two strings), but allows for gaps in the strings. Given a set of edit operations (eg insertions and deletions) and edit costs, the edit distance is given by the minimum sum of the costs along an edit path converting one object into the other. Please consult the Vienna RNA package documentation for a precise definition of string edit distance

<sup>5</sup><http://parallel.hpc.unsw.edu.au/getaegisdist.cgi/getsource/mafold/> version D1

the well-known RNA folding map that illustrated this explanation.

This finding is potentially important. It has been argued that neutral diffusion is an important feature of evolutionary processes allowing efficient search of phenotype space. The sort of competition effects seen here to impede neutral diffusion are characteristic of climax ecosystems. This would imply that disturbed ecosystems will have greater evolvability than climax systems. This “brake” on neutral diffusion being released during times of environmental stress could provide an alternative explanation for the patterns of adaptive radiation seen after mass extinction events.

## Acknowledgments

I would like to thank the *Australian Centre for Advanced Computing and Communications* for a grant of computing time used in this project.

## References

- Adami, C. (1998). *Introduction to Artificial Life*. Springer.
- Barnett, L. (1998). Ruggedness and neutrality — the *NK $\rho$*  family of fitness landscapes. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Artificial Life VI*, pages 18–27, Cambridge, Mass. MIT Press.
- Eldridge, N. (1985). *Time Frames — The Rethinking of Darwinian Evolution and the Theory of Punctuated Equilibria*. Simon and Schuster, New York.
- Huynen, M., Stadler, P. F., and Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proc. Nat. Acad. Sci. USA*, 93:397.
- Kimura, M. (1983). *The Neutral Theory of Molecular Evolution*. Cambridge UP, Cambridge.
- Ray, T. (1991). An approach to the synthesis of life. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, page 371. Addison-Wesley, Reading, Mass.
- Reidys, C., Kopp, S., and Schuster, P. (1997). Evolutionary optimization of biopolymers and sequence structure maps. In Langton, C. and Shimohara, K., editors, *Artificial Life V*, page 379, Cambridge, Mass. MIT Press.
- Schuster, P., Fontana, W., Stadler, P. F., and Hofacker, I. L. (1994). From sequences to shapes and back: A case study in RNA secondary structures. *Proc. Royal Soc. London B*, 255:279–284.
- Standish, R. K. (1997). Embryology in Tierra: A study of a genotype to phenotype map. *Complexity International*,

- Standish, R. K. (1999). Some techniques for the measurement of complexity in Tierra. In Floreano, D., Nicoud, J.-D., and Mondada, F., editors, *Advances in Artificial Life: 5th European Conference, ECAL 99*, volume 1674 of *Lecture Notes in Computer Science*, page 104, Berlin. Springer.
- Standish, R. K. (2003). Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, 3:167.
- Standish, R. K. and Galloway, J. (2002). Visualising Tierra's tree of life using Netmap. In Bilotta, E. et al., editors, *ALife VIII Workshop proceedings*, page 171. <http://alife8.alife.org>.