

Two is better than one: A diploid genotype for neural networks

Raffaele Calabretta^{1,3*}, Riccardo Galbiati², Stefano Nolfi¹, Domenico Parisi¹

¹ *Department of Neural Systems and Artificial Life*

Institute of Psychology, National Research Council

Viale Marx 15, 00137 Rome, Italy

tel.: (+39) 6 86 090 233; fax: (+39) 6 82 47 37

e-mail: rcalabretta@ip.rm.cnr.it

² *Department of Biology, University "Tor Vergata"*

Via della Ricerca scientifica, 00133 Rome, Italy

³ *Centro di Studio per la Chimica del Farmaco, National Research Council*

Piazzale A. Moro 5, 00185 Rome, Italy

Abstract. In nature the genotype of many organisms exhibits diploidy, i.e., it includes two copies of every gene. In this paper we describe the results of simulations comparing the behavior of haploid and diploid populations of ecological neural networks living in both fixed and changing environments. We show that diploid genotypes create more variability in fitness in the population than haploid genotypes and buffer better environmental change; as a consequence, if one wants to obtain good results for both average and peak fitness in a single population one should choose a diploid population with an appropriate mutation rate. Some results of our simulations parallel biological findings.

Key words: *adaptation, diploidy, genetic algorithms, genotype-phenotype mapping, neural networks.*

* Correspondence and proofing.

Two is better than one: A diploid genotype for neural networks

Abstract. In nature the genotype of many organisms exhibits diploidy, i.e., it includes two copies of every gene. In this paper we describe the results of simulations comparing the behavior of haploid and diploid populations of ecological neural networks living in both fixed and changing environments. We show that diploid genotypes create more variability in fitness in the population than haploid genotypes and buffer better environmental change; as a consequence, if one wants to obtain good results for both average and peak fitness in a single population one should choose a diploid population with an appropriate mutation rate. Some results of our simulations parallel biological findings.

1. Introduction

In simulations that apply genetic algorithms to populations of entities the researcher must decide what is encoded in the genotype and how the genotype maps into the phenotype. In most work applying genetic algorithms to populations of neural networks the genotype directly encodes the phenotypical properties of the networks, e.g., the network's connection weights, so that the genotype-to-phenotype mapping is one-to-one. More complex mappings from the genotype to the phenotypical network have been proposed (cf., e.g., [1]) but the variety of genotype-to-phenotype mappings that are found in real organisms is still largely to be explored and analyzed with simulations.

In real organisms genotypes can be either haploid or diploid. A haploid genotype contains a single set of instructions for constructing the phenotype. A diploid genotype contains two distinct sets of instructions, i.e., two sets of chromosomes each of which would be sufficient for constructing the phenotype. Since the two sets of instructions tend to be nonidentical the phenotype is actually constructed by using some instructions from one set and some instructions from the other set. This mapping function results in the expression (*dominance*) and nonexpression (*recessivity*) of particular fragments of the genetic code.

In this paper we report some work that compares the performance of haploid and diploid genotypes in populations of ecological neural networks, i.e., networks controlling the behavior of organisms living in a

(simulated) physical environment [2]. We describe the results of simulations comparing the behavior of haploid and diploid populations in both fixed and changing environments.

2. Simulations

The organism is a simulated mobile robot called Khepera [3,4]. The environment is a rectangular box of 60x35 cm and the robot is rewarded (in the evolutionary sense) for exploring as much as possible of this environment. The environment includes a circular food area of 20 mm of diameter that may be located in different positions in the environment. The robot has a food store that is full when the robot starts exploring the environment and becomes progressively more empty during the exploration. When the food store is completely empty the robot must find the food area. The robot's behavior is controlled by a feed-forward neural network with 10 input units (8 units encoding the activation level of the robot's 8 infrared sensors, 1 unit encoding the current food level of the robot's store, and 1 unit encoding whether the robot is inside or outside the food area). Two output units encode the speed of the robot's two wheels and 2 bias units are connected to these output units.

A genetic algorithm [5] was used to evolve the connection weights of a population of such organisms. An initial population of 100 neural networks was generated by assigning random weights to the 22 connections of each network. The 100 individuals were tested to determine their fitness by placing each of them in a separate copy of the environment. The fitness of an individual was a function of both the amount of environment the individual was able to explore and the ability of the organism to return to the food area when its food store was empty. The 20 individuals that obtained the highest fitness score were allowed to reproduce by generating five copies of their genotype with the addition to some random mutations. The 20x5 new individuals constituted the next generation. The process was repeated for 300 generations.

We used two different genotypes in two different sets of simulations. One genotype was haploid, the other diploid. The haploid genotype included 22 chromosomes (n), one for each of the 22 connection weights of the neural network. Each chromosome was a sequence of 8 bits which coded for a specific value of the corresponding connection weight. The diploid genotype included 22 pairs of chromosomes ($2n$). Each pair of chromosomes coded for two (normally different) values of the same connection weight. Each of the two homologous chromosomes was a sequence of 10 bits. 8 of the 10 bits coded for weight value, as in the haploid genotype (*structural gene*). The remaining 2 bits coded for the dominance/recessivity mechanism (*dominance*

modifier genes; see fig. 1). Notice that in nature dominance can vary in degree, sometimes because of intervening modifiers genes that enhance (or decrease) the expression of another gene (or group of genes) involved in a trait's expression [6].

Conflicts that occurred in decoding the information contained in the structural genes of two homologous chromosomes were resolved by the following mechanism: the first modifier gene was compared with the corresponding modifier gene of the other chromosome and, by using the *dominance rule* (XOR rule: 00 and 11 give 0; 01 and 10 give 1), it was decided which of the two structural genes would be expressed if the second modifier gene was different in the two chromosomes (see fig. 1). If the two chromosomes had an identical second modifier gene, however, the value of the connection weight was the average of the values specified by the structural genes in the two chromosomes (see fig. 1).




Fig. 1. A pair of homologous chromosomes in a diploid genotype. In this case, in both homologous chromosomes the first modifier gene is 1 and therefore, according to the dominance rule, 0 dominates. Since the second modifier gene is 1 in the first chromosome and 0 in the second one, the structural gene of the second chromosome gets expressed. By decoding the two 8 bit-strings and normalizing between -10.0 and +10.0, the left structural gene codes a value of -2.0 while the right one codes a value of +7.6. The actual connection weight value will then be +7.6. If the second modifier gene would have been either 1 or 0 in both chromosomes, the weight value would have been the average of -2.0 and +7.6, that is, +2.8.

3. Results

We have done a number of simulations comparing populations of neural networks with haploid and diploid genotypes. Notice that although diploids have more genes than haploids (the double number of structural genes plus modifiers genes), the resulting neural networks have the same number of weights.

In our simulations we have varied (a) the mutation rate, and (b) whether the population lives in an unchanging environment or in an environment that changes each given number of generations. In this paper we describe the results concerning haploid and diploid populations living either in a fixed environment or in an environment that remains fixed for the first 59 generations and then shifts back and forth between two different forms each 25 generations.

Figure 2 gives the average and peak fitness of haploids (left) and diploids (right) living in the fixed environment for mutation rates of 1%, 2%, and 3%.




Fig. 2. Average and peak fitness of haploid (left) and diploid (right) populations across 300 generations with mutation rates of 1% (black thick curve), 2% (grey curve), and 3% (black thin curve). Average of 6 different replications of the simulation.

Both average and peak fitness of diploids increase linearly with decreasing mutation rate whereas for haploids this is true for average fitness but for peak fitness there is a reversal and the best results are obtained with a mutation rate of 3%. Haploids obtain better results than diploids in average fitness for all mutation rates but the best peak fitness is reached by diploids with mutation rate of 1%.

Figure 3 compares average and peak fitness of diploids with mutation rate of 1% and of haploids with mutation rate of 1% (left) and 3% (right) in the changing environment. These conditions were chosen because, in a fixed environment, diploids with 1% mutation rate obtain the best results for both average and peak fitness of all diploids while among haploids the best average fitness is reached with 1% mutation rate and the best peak fitness with 3% mutation rate.




Fig. 3. Average and peak fitness in a diploid population with 1% mutation rate (black curves) and in a haploid population (grey curves) with 1% (left) and 3% (right) mutation rate. Average of 6 different replications of the simulation.

In analyzing the results of figure 3, one must consider that in the fixed environment condition one can compare two different populations only with respect to rapidity of fitness increase and to fitness level reached after 300 generations (when fitness is in a more or less steady state). If the environment changes one can also compare the two populations with respect to the amount of fitness decrease observed in the first generation after the environment has changed.

Figure 3 shows that if one wants to obtain good results for both average and peak fitness with a single population one should choose a diploid population with a mutation rate of 1%. When compared with haploids with 3% mutation rate this population of diploids obtains better results for both average and peak fitness (right). Not only the diploids have more rapid fitness increase and higher terminal fitness but their fitness appears to be less negatively affected by environmental change. When compared with haploids with 1% mutation rate (left) diploids with 1% mutation rate obtain slightly less good results for average fitness but only with respect to level of terminal fitness. On the other hand, the diploids obtain better peak fitness than the haploids. No single haploid population is able to reach these results.

4. Discussion

Let us consider the results with fixed environment first. Generally mutations tend to have larger effects on diploids than on haploids because diploids have modifier genes that are lacking in haploids. If a mutation affects a modifier gene its effects can be very large and in any case much larger than those of any mutation

affecting structural genes. Since most effects of mutations are deleterious more individuals are negatively affected by mutations in diploids than in haploids. Therefore, we find that average fitness is lower in diploids than in haploids with any given mutation rate.

On the other hand, since in diploids mutations tend to have larger effects it is more likely that some rare advantageous mutation results in some particularly good individual in diploids than in haploids. In other words, diploidy tends to push average fitness down but it may also cause peak fitness to go up. Hence, diploids tend to have lower average fitness but higher peak fitness than haploids. More generally the frequency distribution of fitness values tends to be bimodal in diploids and unimodal in haploids (cf. figure 4).




Fig. 4. Frequency distribution of individuals with various fitness values across 300 generations (sampled each 5 generations) in an unchanging environment for haploids (left graph) and diploids (right graph) with mutation rate of 1% in two representative simulations.

However, these results are obtained only for mutations rates of 1% and 2%. A mutation rate of 3% creates too much disruption in the diploid population which therefore turns out to have both lower average and peak fitness than the haploid population.

These results hold for both fixed and changing environments. However, in the changing environment diploids exhibit another feature that characterizes them with respect to haploids, i.e., their capacity to keep a genetic "memory" of the past that can be useful when the population must re-adapt to an environment to which it has already adapted in the past. Haploids have all their genes expressed and therefore their entire genetic endowment becomes adapted to the current environment. When the environment changes, the negative effects on fitness are felt more strongly by haploids than by diploids whose genetic "memory", recorded in their non-expressed genes, allows them to be less negatively affected by environmental changes.

In their theoretical and experimental analysis of diploidy and dominance applied to a 0-1 knapsack problem, Goldberg and Smith [7] have shown that diploid genotypes perform in a temporally varying environment better than haploid ones because "diploidy embodies a form of temporal memory that is distributed across population". Various studies have focused on the use of diploid genotypes and dominance operators in artificial genetic search (e.g., [8] and [9]) but, to our knowledge, the present work is the first one comparing haploid and diploid genotypes for the optimization of neural networks. The study of diploidy in this type of optimization problem is a hard one because due to epistatic interactions the quantity and quality of effects on a neural network's fitness of a single weight value depend on the values of many other weights. As a consequence, it is less obvious that different alleles may encode solutions suited for different environments.

From a biological point of view, the experimental results of Paquin and Adams [10] on yeast populations are those more usefully comparable with our simulation results. In some simulations not reported here in which the environment shifts from A to B and back every other generation, diploid populations with 1% mutation rate show very little fluctuation in fitness value compared with haploids [11]. In fact, they seem to have found a good solution for both environments. This last result supports the conclusions of Paquin and Adams [10] on the role of diploidy in short-term adaptation to new environments in asexual organisms. More generally, we think that a deeper theoretical analysis of our approach can provide useful insights on the significance and evolution of diploidy (cf. [11]).

References

1. S. Nolfi and D. Parisi, "Genotypes for Neural Networks", in M. A. Arbib (ed.) *The Handbook of Brain Theory and Neural Networks*, pp. 431-434, Bradford Books: MIT Press, Cambridge, 1995.
2. D. Parisi, F. Cecconi and S. Nolfi, "Econets: Neural Networks that Learn in an Environment", *Network*, Vol. 1, pp. 149-168, 1990.
3. F. Mondada, E. Franzi and P. Ienne, "Mobile Robot Miniaturisation: a Tool for investigation in control algorithms", in *Proc. of the 3rd Int. Symposium on Experimental Robotics*, Kyoto, 1993.
4. S. Nolfi, D. Floreano, O. Miglino and F. Mondada, "How to Evolve Autonomous Robots: Different Approaches in evolutionary robotics", in R. Brooks and P. Maes (eds) *Artificial Life IV: Proc. of 4th Int. Workshop on the Synthesis and Simulation of Living Systems*, pp. 190-197, Cambridge, MIT Press, 1994.

5. J. H. Holland, *Adaptation in Natural and Artificial Systems*, The University of Michigan Press: Ann Arbor, 1975.
6. B. Wallace, *Basic Population Genetics*, Columbia University Press: New York, 1981.
7. R. E. Smith and D. E. Goldberg, "Diploidy and Dominance in Artificial Genetic Search", *Complex Systems*, Vol. 6, pp. 251-285, 1992.
8. E. Collingwood, D. Corne and P. Ross, "Useful Diversity via Multiploidy", in *Proc. of IEEE 3rd Int. Conf. on Evolutionary Computation*, Nagoya, 1996.
9. C. Fonteix, F. Bicking, E. Perrin and I. Marc, "Haploid and Diploid Algorithms, a New Approach for Global Optimization: Compared Performances", *Int. J. Systems Sci.*, Vol. 26, pp. 1919-1933, 1995.
10. C. Paquin and J. Adams, "Frequency of Adaptive Mutations is Higher in Evolving Diploid than Haploid Yeast Population", *Nature*, Vol. 230, pp. 495-500, 1983.
11. R. Calabretta, R. Galbiati, S. Nolfi. and D. Parisi, "Haploid and Diploid Genotypes for Neural Networks", Institute of Psychology, C.N.R., Rome, Technical Report NSAL-96005, 1996.