Investigating the role of diploidy in simulated populations of evolving individuals

Raffaele Calabretta[1,3], Riccardo Galbiati[2], Stefano Nolfi[1] and Domenico Parisi[1]

[1] Department of Neural Systems and Artificial Life
Institute of Psychology, National Research Council
Viale Marx 15, 00137 Rome, Italy
e-mail: rcalabretta@ip.rm.cnr.it

[2] Department of Biology, University "Tor Vergata"
Via della Ricerca scientifica, 00133 Rome, Italy

[3] Department of Ecology and Evolutionary Biology, Yale University
165 Prospect Street, 06511 New Haven, CT, U.S.A.

Abstract. In most work applying genetic algorithms to populations of neural networks there is no real distinction between genotype and phenotype. In nature both the information contained in the genotype and the mapping of the genetic information into the phenotype are usually much more complex. The genotypes of many organisms exhibit diploidy, i.e., they include two copies of each gene: if the two copies are not identical in their sequences and therefore have a functional difference in their products (usually proteins), the expressed phenotypic feature is termed the dominant one, the other one recessive (not expressed). In this paper we review the literature on the use of diploidy and dominance operators in genetic algorithms; we present the new results we obtained with our own simulations in changing environments; finally, we discuss some results of our simulations that parallel biological findings.

1 Genotypes for neural networks

Genetic algorithms are computational models of evolution that may be applied to populations of neural networks to study the evolution of organisms whose behavior is controlled by a neural network. Imagine a population of such organisms living in an environment and reproducing as a function of some performance criterion. The initial population of neural networks is randomly generated. Hence, each individual organism will be assigned a neural network that is different from the neural network of other individuals and, as a consequence, will tend to behave differently from any other individual. The individuals that behave more efficiently according to the performance criterion will reproduce while less fit individuals are more likely to die without leaving offspring. Reproduction consists in generating one or more copies of the reproducing individual's neural network (we are assuming non-sexual reproduction) with the addition of some random changes to some of the network's traits (genetic mutations). Hence, an offspring's behavior will tend to be similar but not identical to its parent's behavior. Genetic mutations will result in most cases in offspring that perform less well than their parents but in some rare cases an offspring will outperform its parent. It is these luckier individuals that will tend to reproduce rather than their less lucky siblings. Selective reproduction and the constant addition of variability to the "genetic pool" will cause an increase in the population's level of performance across a certain number of generations.

In most work applying genetic algorithms to populations of neural networks there is no real distinction between genotype and phenotype. The inherited genetic information maps one to one into the phenotypic neural network. For example, in populations with fixed network architecture the genetically inherited trait tends to be the matrix of connection weights. The inherited genotype directly encodes the matrix of weights and the genetic mutations directly change the value of some of the weights. Therefore, the genotype and the phenotype are virtually identical. In nature both the information contained in the genotype and the mapping of the genetic information into the phenotype are usually much more complex. The mapping from genetic to phenotypical traits is not one to one, but one single genetic trait can enter into the determination of many different phenotypical traits

(*pleiotropy*) and, vice versa, one phenotypical trait can be determined by the concurrent action of many different genetic traits (*poligeny*). Furthermore, the genotype contain not only information that directly maps into the phenotype but also higher order information that regulates the mapping. Finally, it should be considered the external environment is an additional factor that, by interacting with the genetic information, determines the phenotype (for an interesting discussion on evolution of genotype-phenotype mapping see Wagner and Altenberg, 1996).
Several researchers have already attempted to study more complex and biologically plausible genotype/phenotype mappings (Cangelosi *et al.*, 1994; Nolfi & Parisi, 1995; Dellaert & Beer, 1994), but the variety and complexity of genotype-to-phenotype mappings found in real organisms is still largely to be explored and analyzed with simulations.
In particular, the genotypes of many living beings exhibit diploidy, i.e., they include two copies of each gene. In our previous short paper (Calabretta *et al.*, 1996) we described the results of simulations comparing the behavior of haploid and diploid populations of ecological neural networks living in both fixed and changing environments.
In this paper we review the literature on the use of diploidy and dominance operators in genetic algorithms. We present the new results we obtained with our own simulations in changing environments. Finally, we discuss the insights this approach can provide to understand the conditions in which diploidy can enhance the adaptation power of asexual organisms.

2 Haploidy and diploidy

2.1 Artificial Life perspective

In a section of his well known book about Genetic Algorithms (GAs), Goldberg (1989) wondered whether diploid genotype and dominance operators can be useful in artificial genetic search.
Several studies have focused on the use of diploid genotypes and dominance operators in genetic algorithms. Historically the first attempts date back to the beginning of the seventies. Hollstien (1971) introduced a model with diploidy and an evolving dominance mechanism based on a triallelic scheme and, some years later, Holland (1975) discussed and analyzed the steady-state performance of this model. According to Goldberg (1989) "Hollstien-Holland triallelic scheme is the clearest, simplest scheme suggested for artificial genetic search thus far, combining both dominance map an allele information at a single position. With this scheme the more effective allele becomes dominant , thereby shielding the recessive."
Smith and Goldberg (1992) stressed the role of diploidy and dominance as abeyance structures and mechanisms. In their theoretical and experimental analysis of diploidy and dominance applied to a 0-1 knapsack problem (Syslo *et al.*, 1983), which belongs to a class of common but difficult problems in operations research, they demonstrated that diploid GAs perform in temporally varying environment better than haploid GAs because " ... diploidy embodies a form of temporal memory that is distributed across population" and that "... an adaptive dominance map is necessary to effectively exploit the advantages of diploidy."
Fonteix *et al.* (1995) compared, with regard to the convergence time, haploid and diploid algorithms on several complex optimization problems by measuring the number of generations needed to reach the solution. They stressed that for simple problems the performances of the two algorithms were similar, while for more complex problems the diploid algorithm needed a fewer generations.
Collingwood *et al.* (1996) stressed the usefulness of multiploid GA “in cases where attractive suboptima are profoundly Hamming distant from the true optimum, thus requiring a GA to recover substantial lost material in order to recover from suboptima.”
To our knowledge, the present work is the first simulative attempt to compare haploid and diploid genotypes in asexual populations. This, by excluding another complicated factor, may prove important to understand the role of diploidy in adaptive individuals. Moreover, in our simulation we present a genotype-to-phenotype mapping more realistic than those used in the works described above in which the genotype codifies for the nervous system of an organism interacting with an external environment.

2.1 Biological perspective

Organisms have a genotype which can be either haploid like bacteria or diploid like most animals and plants. Ploidy means the number of genome copies, that is the complete set of genetic information; so, haploid means one copy, diploid two copies, polyploid more than two copies (like in some plants, or animal tissues).
Genes are encoded by sequences of four different chemicals called nucleotides (the well-known DNA double-helix) and in diploid organisms the two copies of the same gene (alleles) are placed in different but corresponding members of a chromosome pair (e.g., in humans there are 22 pairs, plus one of sexual chromosomes). If the two copies are not identical in their sequences and from this stems a functional difference of their products (usually proteins**)**, the expressed phenotypic feature is termed the dominant one, the other one recessive (not expressed). Dominance can be complete (e.g. brown vs. blue eye colour) or incomplete when the recessive feature is partially expressed (e.g. colour of some hybrid flowers). Co-dominance (i.e. the expression of both features) can also occur, (e.g. AB blood group, Strickberger, 1976). Therefore, in some circumstances, dominance can vary in degree, because of intervening modifiers genes that enhance (or inhibit) the expression of

another gene (or group of genes) involved in a trait's expression (Wallace, 1981).
Diploids are believed to adapt better and faster than haploids for several reasons: (a) diploids can mask the effect of deleterious mutations which usually affects the recessive features of a trait; (b) overdominance (i.e. a positive interaction between different alleles in the expression of a trait) may improve the adaptability of evolving individuals; (c) a larger occurrence of favourable and initially partial dominant mutations (Crow & Kimura, 1965; Paquin & Adams, 1983; Kondrashov & Crow, 1991; Perrot et al., 1991). On the other hand, because diploids are subjected to a larger number of mutations with respect to haploids, a long-term reduction in fitness should be expected unless dominance is complete or strong epistatic effects are present (i.e. decoupling of phenotypic expression relative to its genetic background; Kondrashov & Crow, 1991) (Otto & Goldstein, 1992; Goldstein, 1992; Orr & Otto, 1994).
Several theoretical investigations have been conducted to try to explain the importance of diploidy and how it could have evolved. In fact, the presence of diploidy in many complex organisms (such as in the commonly but erroneously named “higher” plants and animals) and in other groups (such as fungi) suggest that probably other factors, in addition to those outlined above, are involved.
For example, Buss claimed that diploidy can protect individuals from somatic mutations and that, by causing an increase in the cell size, can allow a more rapid tissue growth (Buss, 1987). In this paper, however, we will focus on an aspect of ploidy to which Paquin and Adams (1983) refer in an experimental work on haploid and diploid yeast strains, in which they conclude: "... the rate of *(adaptive)* mutation ... may be critical in determining short term adaptation to new environments for asexual organisms which cannot rely on recombination to generate variation in fitness."

3 Simulations

3.1 The task, the robot and the environment

We ran a set of simulations in which the task of the evolving populations is to explore an environment and return, time to time, to a "food" area where individuals can reintegrate the energy consumed during the exploration. The organism is a miniature mobile robot called Khepera, developed at E.P.F.L. in Lausanne (Mondada *et al.*, 1993). Khepera has a circular shape with a diameter of 55 mm., a height of 30 mm., and a weight of 70 g. The robot is supported by two wheels that allow it to move in various directions by separately regulating the speed of each wheel. The robot is also provided with eight infra-red proximity sensors positioned on the periphery of its body with six sensors on the front side and two sensors on the back. The infrared sensors allow the robot to detect obstacles to a distance of about 4 cm.
The environment is a rectangular box of 6Ox35 cm and contains a circular food area of 2O mm diameter located in a randomly selected position within the box.
The robot has a food store that is full when it starts exploring the environment and becomes progressively more empty during the exploration. When the food store is completely empty, the robot must reach the food area and remain there until the food store is full again. Then it can start a new exploration of the environment. The robot must explore the environment efficiently in the sense that it must visit as much of the environment as possible. In doing so, the robot must develop an ability to avoid hitting the walls, otherwise it would get stuck on the walls themselves.
A simulator of both the robot and the environment was developed by recording samples of the sensory patterns that the real robot perceives in the real environment (Nolfi et al., 1994).
The robot's behavior was controlled by a feed-forward neural network. The network included 10 input units (8 units encoding the activation level of the 8 infrared sensors, 1 unit encoding the current energy level of the robot, and 1 unit encoding whether the robot is inside or outside the food area), 2 output units encoding the speed of the two corresponding robot's wheels. Note that the robot can sense the food area only when it is over the food area itself. Therefore, when the robot must reintegrate its reserve of energy it must find the food area by exploring the environment.
A genetic algorithm (Holland, 1975; Goldberg, 1989; Mitchell & Forest, 1994) 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 weights of each network (20 weights connecting the 10 input units to the 2 output units and 2 bias weights for each output units). The 100 individuals were tested to determine their fitness by placing each of them in a separate copy of the environment. Each individual was placed in the box with a randomly selected orientation and it was allowed to move for 2,000 cycles each corresponding to 100 ms of real time. This process was repeated three times (epochs) for a total of 6,000 cycles. The environment was ideally divided up into cells of 2x2 cm and individuals were scored for the total number of cells visited for the first time during each epoch being the robot’s energy above 0.0 (cells visited when the energy level was below such a threshold did not produce an increase in fitness).
The 20 individuals that obtained the highest fitness score were allowed to reproduce by generating five copies of their genotype with the addition of random mutations. The 2Ox5 new individuals constituted the next generation that was tested exactly like the first one. The process was continued for 300 generations.

3.2 Haploid and diploid genotypes

We used two different types of genetic coding for our neural networks and we ran two different sets of simulations. The first type of genetic encoding was haploid, the second diploid.
The haploid genotype included 22 chromosomes (*n*), one for each of the 22 connection weights of the neural network. Each chromosome is a sequence of 8 bits (0 or 1) which coded for a specific value of the corresponding connection weight. Normal binary coding was used to translate the 8 bits to one weight value between -10.0 and + 10.0.
The diploid genotype included 22 pairs of chromosomes (*2n*) (see Figure 1). Each pair of chromosomes coded for two possibly different values of the corresponding connection weight. In the case of diploid genotypes, each of the two homologous chromosomes consisted of a sequence of 10 bits: 8 coding for the corresponding weight value chromosome (*structural genes)* and the remaining 2 for the dominance/recessivity mechanism (*dominance modifier genes).*
The dominance/recessivity mechanism was implemented in the following way: the first modifier gene was compared with the corresponding modifier gene of the corresponding chromosome and an XOR rule was used to decide which of the two structural genes would be expressed (00 and 11 imply that the chromosome with the second modifier gene equal to 0 will be expressed while 01 and 10 imply that the chromosome with the second modifier gene equal to 1 will be expressed). However, if the two chromosomes had an identical second modifier gene**,** the two homologous chromosomes were considered as *co-dominants* and the average of the two values specified by the two structural genes is used (see Figure 1).

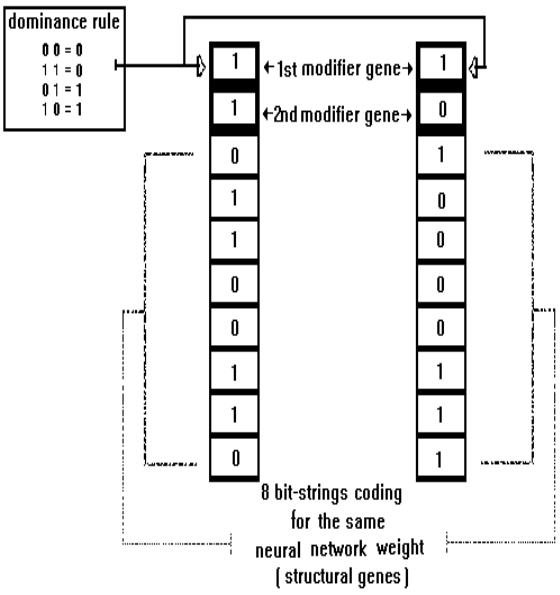

Fig. 1. A pair of homologous chromosomes in a diploid genotype. In this case, the first modifier genes of both homologous chromosomes is 1 and therefore, according to the XOR dominance rule, the second modifier gene 0 dominates. Since the second modifier gene is 1 in the first chromosome and 0 in the second one, the structural gene of the second chromosome is expressed (dominant). 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 had 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.

4 Experiments and Results

We will present several sets of simulations in which we will compare diploid and haploid individuals with different environmental conditions and with different mutation rates.
We ran sets of simulations in which the environment is stable or changing during the evolutionary process. In the first case the position of the food area was 175 on the x coordinate and 175 on the y coordinate and it remained the same for the entire course of evolution (300 generations). In the second case the position of the food area remained the same (175/175) for the first 59 generations; and then started to alternate between two different positions (175/175 and 500/100). In the case of the changing environment we also analyzed two different cases: the case in which the position of the food area changed each 25 generations and the case it changed each generation.
For each condition described above we investigated the performance obtained with different mutation rates (1%, 2%, or 3% of the bits of the genotype randomly selected were replaced with a new randomly selected value). In particular, we were interested in examining the consequences of different mutation rates for haploid and diploid individuals with reference at distribution of the fitness among the individuals of the population.
For each condition we ran 6 experiments starting with different randomly assigned genotypes. Each simulation lasted 300 generations.

4.1 Unchanging environment

Figure 2 gives the average and peak fitness (respectively, the average fitness of the population and the fitness of the best individual of the population) of haploids (top) and diploids (bottom) living in the fixed environment for mutation rates of 1%, 2%, and 3%.
Both average and peak fitness of diploids and the average fitness of haploids increase linearly with a decreased mutation rate. However the haploids reach the highest peak fitness levels with a mutation rate of 3%. Haploids obtain better results than diploids in average fitness for all mutation rates (the best average fitness is reached by haploids with mutation rate of 1%). However, diploids

overcome haploids in the peak fitness (the best peak fitness is reached by diploids with mutation rate of 1%). This is shown more clearly in Figure 3 (top graph), which shows the average and peak fitness throughout 300 generations for haploid and diploid populations with a mutation rate of 1%. The average fitness is lower but the fitness of the best individuals is higher in the diploid populations than the haploid populations. In other words, in the diploid populations there is more distance between the average and the fitness of the best individual than in haploids. There is more variability in fitness among individuals of the same generations in diploids than in haploids.

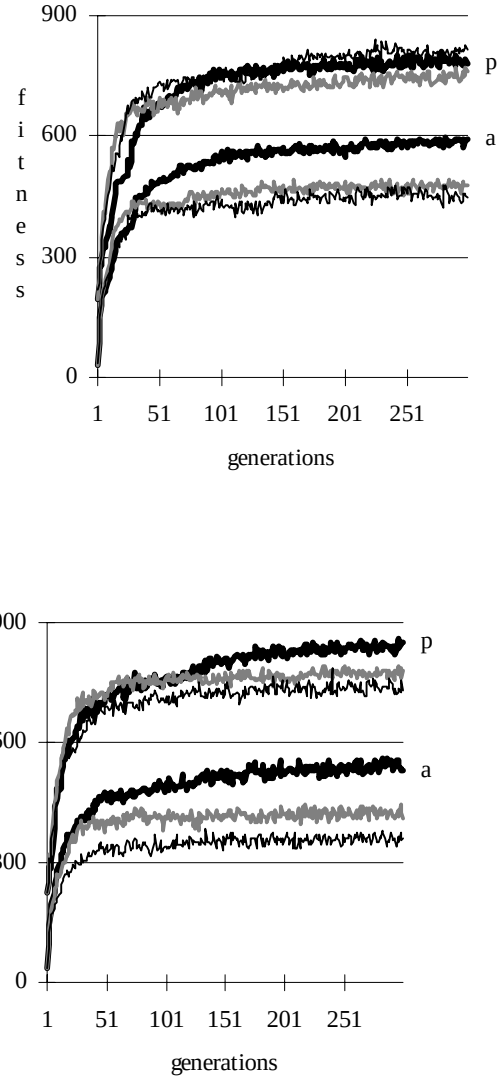

Fig. 2. Average and peak fitness of haploid (top) and diploid (bottom) populations across 300 generations with mutation rates of 1% (thick black curve), 2% (gray curve), and 3% (thin black curve), living in a fixed environment. Average of 6 different replications of the simulation.

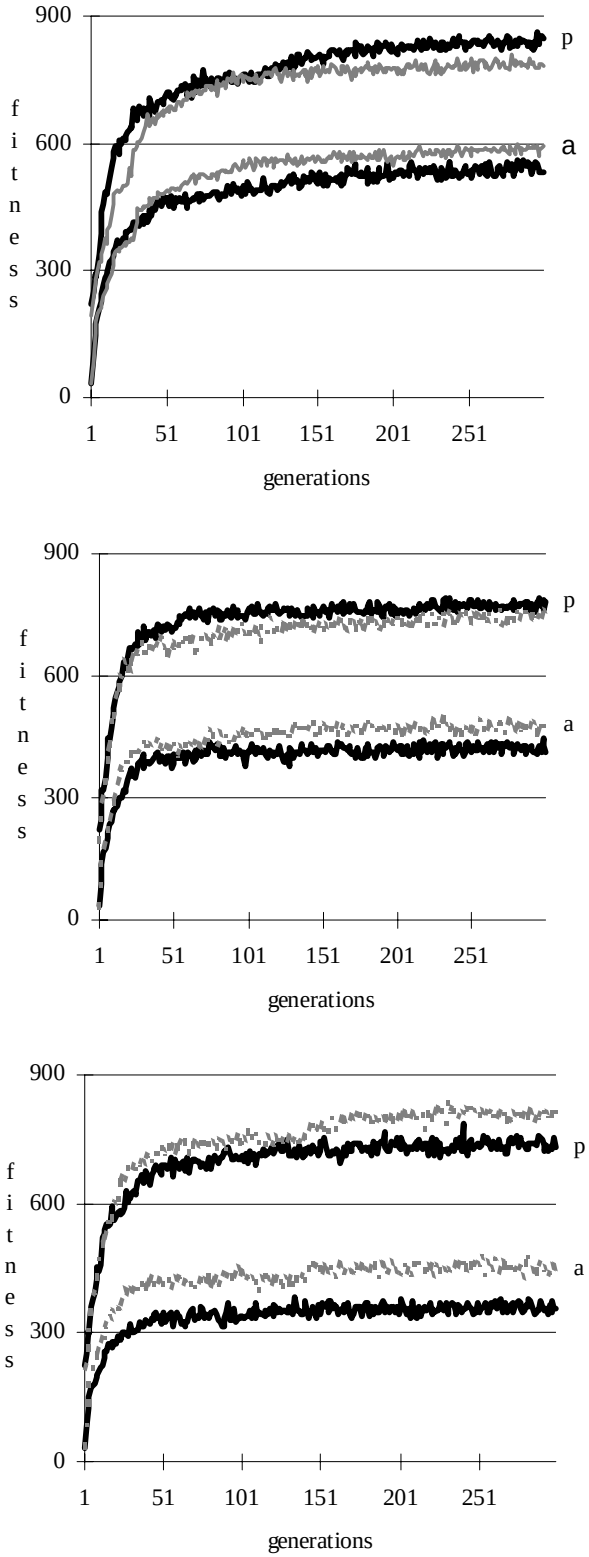

Fig. 3. Average fitness and fitness of the best individual throughout 300 generations in an unchanging environment for diploids (black curve) and haploids population (gray curve) with a mutation rate of 1% (top graph), 2% (middle

graph) and 3% (bottom graph). Average of 6 different replications of the same simulation.

This difference in fitness variability within individuals of the same population can be directly observed by comparing the frequency distribution of different fitness values in each generation for haploid (top graph) and diploid (bottom graph) individuals (see Figure 4).
In haploid populations most individuals have an average level fitness and few individuals have a much higher level of fitness. On the other hand, diploid populations have about half of the population with very low level fitness but also tend to include individuals that have average level, good and very good fitness values. Moreover, as we said above, the best diploid individuals overcome the best haploid individuals in fitness level.

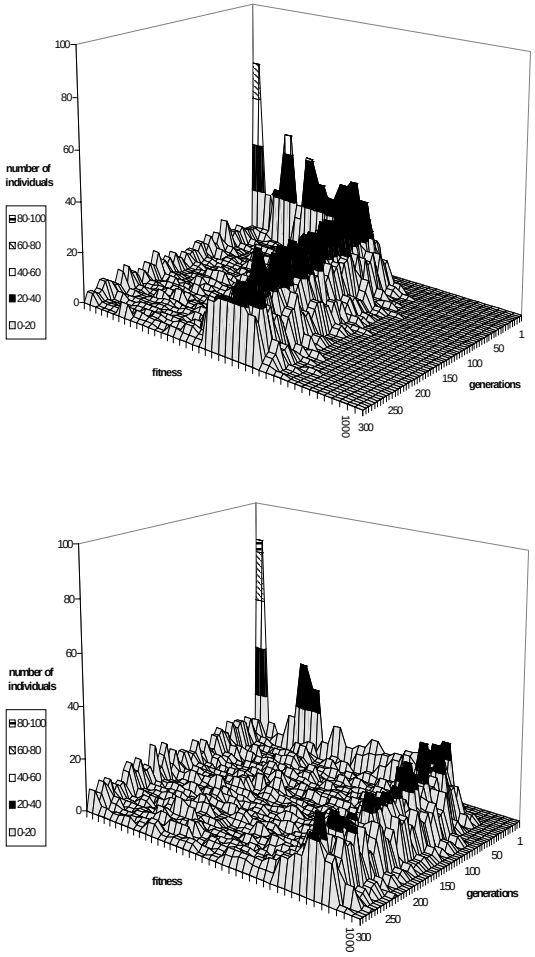

Fig. 4. Frequency distribution of individuals with different fitness values throughout 300 generations (sampled each 5 generations). Fitness is divided into 21 fitness ranges, which go from zero to 1000. Data from two representative simulations made in an unchanging environment for haploids (top graph) and diploids (bottom graph) with mutation rate of 1%.

These differences in the distribution of fitness values in the two populations appear to be a result of the different effects of mutations on haploid and diploid genotypes. On one side, the same rate of genetic mutations (1%) has a more disruptive effect in diploid than in haploid populations. Consider that the 100 individuals of each generation are the offspring of the 20 individuals of the preceding generation that have the highest fitness. Genetic mutations can be said to have a disruptive effect in so far as they cause the offspring of these 20 individuals with high fitness to have a very low fitness. We see that this happens to a greater extent in the diploid than in the haploid populations. On the other hand, diploid populations appear to be more efficient in exploring the fitness landscape in that in each generation there are a few individuals that have a higher fitness than the corresponding best individuals of the haploid populations. In other words, in a diploid population it is more probable that a mutation results in an offspring that is more fit than its parent. Taken together, these two differences in the effects of mutations in haploid and diploid populations can be interpreted as a greater addition of variability in the genetic pool of mutations operating in diploid rather than in haploid genotypes. 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).
However, these results are obtained only for mutations rates of 1% and 2% (cf. the top and middle graphs of figure 3). 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 (cf. the bottom graph of figure 3).
The different effects of genetic mutations in haploid and diploid populations can be explained if we consider how haploid and diploid genotypes can be affected by mutations. In a haploid genotype a single mutation, i.e., a change in bit value, can only affect the structural information contained in a gene. In our case this means that the value of a connection weight can be changed by a mutation but the change can only be a more or less great divergence (depending on the position of the bit that is changed in the sequence of 8 bits) of the new weight value from the old weight value. Furthermore, it is possible that a positive (excitatory) weight value is changed by a mutation to a negative (inhibitory) value, or vice versa, but this cannot but be a rare event because it can only happen if the old weight value is near 0 and/or the change is sufficiently great.
Consider now what the effect of a single mutation can be in the case of a diploid genotype. We must distinguish between the case in which the mutation affects the structural portion of a gene and the case in which the

mutation affects the regulatory portion of the gene, i.e., the two bits encoding the dominance/recessivity mechanism. If the mutation affects the structural portion of a gene, the effects are similar to those we have just analyzed for the haploid genotype. The only difference is that the mutated gene in the diploid genotype can be non-expressed and in this case there is no visible effect of the mutation on the phenotype. However, the effect of the mutation is not lost because the mutated gene remains as part of the genotype and it can be become expressed in some descendant of the current individual. (See below where we will discuss the results of the simulations with changing environments). But if a mutation operating in a diploid genotype affects the regulatory portion of a gene instead of its structural portion, its effects can be much greater than the effects of mutations in haploid genotypes. The mutation can change the decisions taken by the dominance/recessivity mechanism of the gene and, as a consequence, the phenotypic trait of the offspring controlled by the gene can now be determined by the homologous gene which in the parent was non-expressed. This can lead to changes in the phenotype that are more radical than those that can occur in a haploid genotype. In our case the new value of a connection weight tends to be uncorrelated with the old weight and, furthermore, there is a much higher probability that an excitatory connection is changed to a inhibitory connection, or vice versa.

4.2 Changing environments

The different impact of adding more or less variability on average vs. best fitness is even clearer when the population must adapt to an environment that changes periodically after a certain number of generations. Adding more variability to the population's genetic pool results in a higher fitness for the population's best individuals but to lower average fitness.
Figure 5 shows the average and peak fitness of haploid and diploid individuals in changing environments (i.e. environments in which the position of the food area changes each 25 generations) with mutation rates of 1% and 3% (see also Calabretta et al, 1996). These mutation rates were chosen because, in a fixed environment, diploids with a 1% mutation rate obtain the best results for both average and peak fitness while among haploids the best average fitness is reached with a 1% mutation rate and the best peak fitness with a 3% mutation rate.
In analyzing the results of Figure 5, it is important to notice that while in a fixed environment one can compare two different populations only with respect to level of fitness reached after a certain number of generations and to the time need to reach such level, in a changing environment, one can compare two populations also with respect to the amount of fitness decrease observed after a change in the environment.

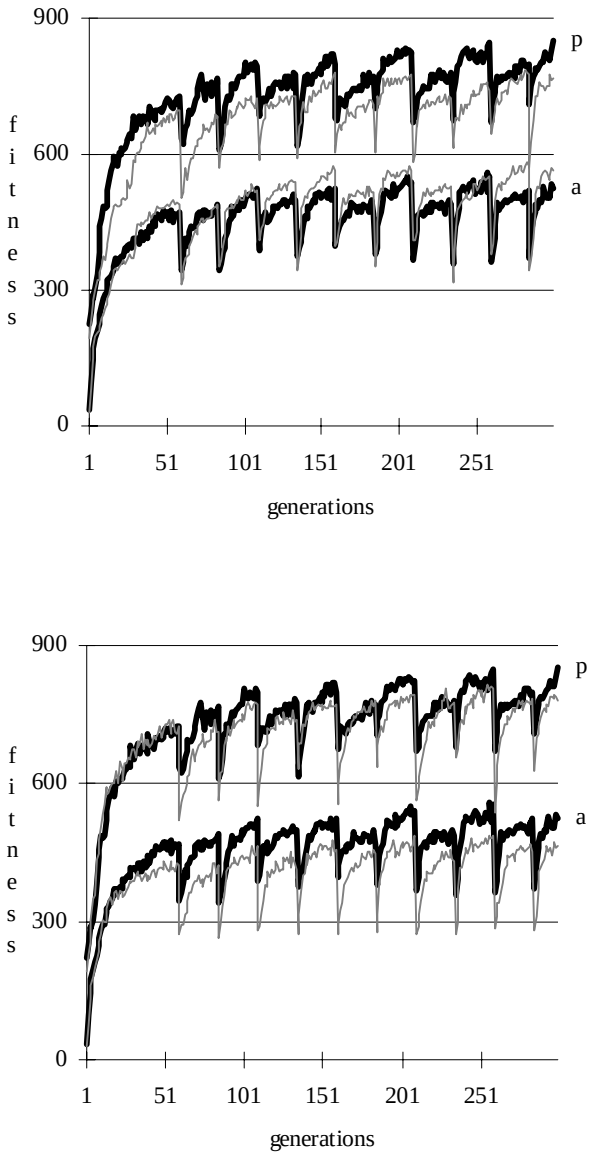

Fig. 5. Average and peak fitness in a diploid population with 1% mutation rate (black curves in the top and bottom graphs) and in a haploid population (gray curves) with 1% (top graph) and 3% (bottom graph) mutation rates. Individuals were placed in an environment that after the first 59 generations started to change every 25 generations. Average of 6 different replications of the simulation.

Figure 5 shows that the best results overall are obtained with diploids and a mutation rate of 1%. In fact, in performance they overcome haploids with a mutation rate of 3% both from the average and peak performance point of view. Diploids not only have a more rapid fitness increase and a higher level of fitness before the next change in the environment, but their performance appears to decrease less when the environmental conditions suddenly change.

When compared with haploids with a mutation rate of 1% (top graph), diploids with the same mutation rate obtain a slightly lower average fitness level only at the end of the period in which the environment does not change. However, they obtain better peak fitness than the haploids. Similar results are obtained when individuals are placed in environments that change more or less frequently (i.e. each 10 or 50 generations).
On the other hand, when the environment changes each generation we observed a different picture. Figure 6 compares average and peak fitness of diploids with a mutation rate of 1% and of haploids with a mutation rate of 1% (top graph), 2% (middle graph) and 3% (bottom graph): diploids with mutation rate of 1% are the only ones able to tolerate continuous environmental fluctuations by presenting only small fluctuations in both their average and peak fitness.

5 Discussion

From a biological point of view, the theoretical findings of Orr and Otto (1994) relative to the rate of adaptation in asexual haploids and diploids and the experimental results of Paquin and Adams (1983) on yeast populations can more easily be compared with our simulation results.
Orr and Otto (1994) considered the rate at which favourable (adaptive) mutations appear and spread in populations of asexual haploids and diploids. They found that the rate of incorporation of favourable mutations in diploids depends on the dominance of advantageous mutations (h) and on the number of favourable mutations per generation (νN), where N is the population size and ν is the rate at which favourable mutations appear per haploid genome. Because diploidy doubles the rate of occurrence of favourable mutations, diploids are expected to adapt faster than haploids but, according to Orr and Otto (1994), only when there is a high level of dominance and a small product νN (to avoid accidental loss of favorable alleles and a consequent reduction in selective advantage in heterozygotes and to allow fixation of favorable mutations respectively).
We have not explicitly investigated the rate of adaptation or the proportion and the rate of occurrence of deleterious vs. favorable mutations. However, we have analyzed, in simulations, the performance of haploid and diploid asexual individuals with populations of constant size in a changing environment. These individuals also included a mechanism that allows dominance shift (through the mutation of modifier genes) and a mapping genotype/phenotype with strong epistatic interactions (due to the non-linear interactions between the neural network's weights)

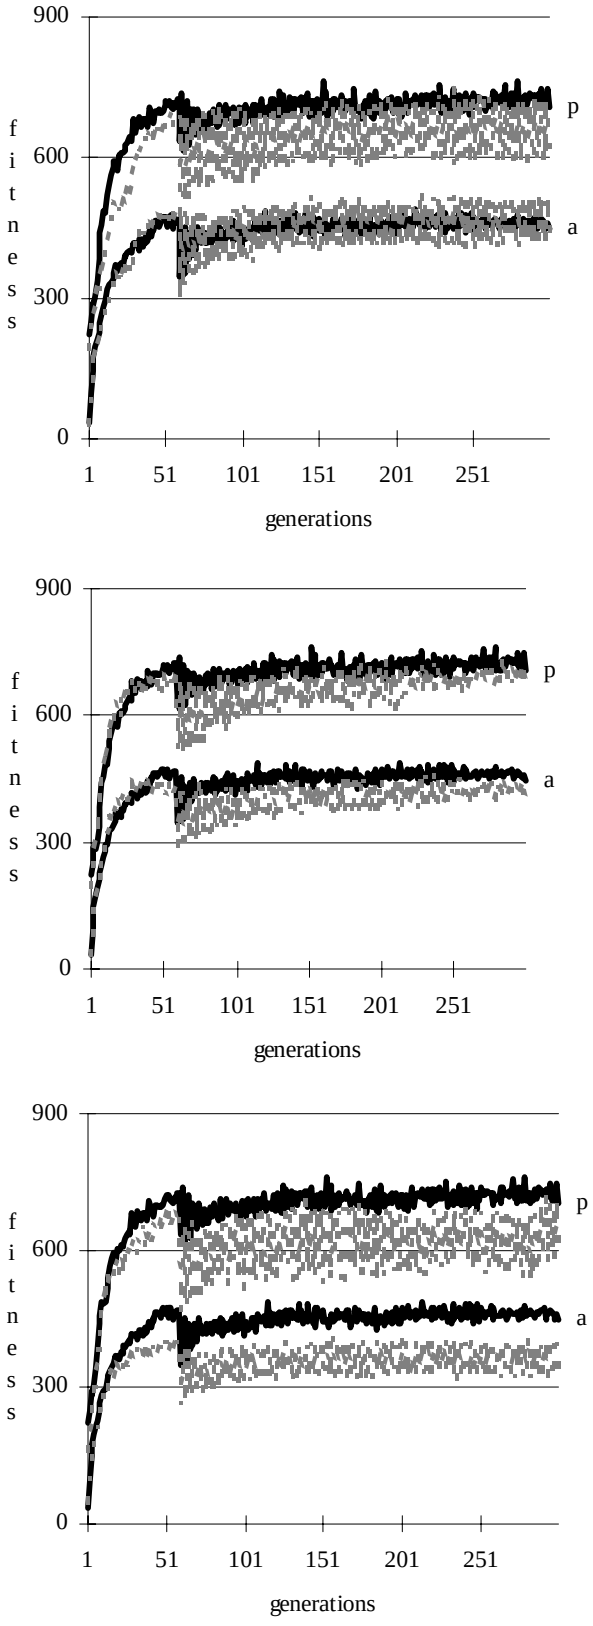

Fig. 6. Average and peak fitness in a diploid population with a mutation rate of 1% (black curve in the top, middle and bottom graph) and in haploid population (gray curve)

for mutation rates of 1% (top graph), 2% (middle graph), and 3% (bottom graph). Individuals were placed in an environment that after 59 generations started to change each generation. Average results of 6 different replications of the simulation.

Nevertheless, our results seems to confirm the hypothesis of Orr and Otto (1994) if we consider the distribution of different fitness value among the individuals of the population and the tolerance to environmental changes. In fact diploids with a mutation rate of 1% overcome in performance (especially in average performance) diploids with higher mutation rates, no matter what the rate of world change is (indeed, similar results are obtained with haploids). Moreover, in changing environments diploids with a mutation rate of 1% overcome in performance haploids (no matter what their mutation rate is), except for the average performance of the haploids with a mutation rate of 1%. The higher performance values in diploids with a 1% mutation rate can be explained with the effect of dominance and epistasis that allows, through mutations, non-expressed genetic information to be explored more freely to allow, occasionally, the development of a new adaptive complexes that may later be extremely useful especially after an environmental change. On the other hand, in haploid organisms only epistatic effects are present and a great deal of old and new (mutated) information is directly expressed in phenotype and therefore subjected to selection.
Another possible explanation of the fact that diploids (with a mutation rate of 1%) overcome haploids in changing environments may be that diploids, by exploiting the dominance mechanism, are able to preserve traits that are temporarily not useful, because of an environmental change, but that can be useful again when the environment changes back again. In our simulations diploids with a mutation rate of 1% present fluctuations in performances comparable with those of haploids when the environment changes each 25 generations. However, when the environment change more often (each generation) diploids with a mutation rate of 1% show much lower fluctuations in their fitness values than haploids. The fact that they seem to be able to find a good solution to both environments may support the claims of Paquin and Adams (1983) on the role of diploidy in short-term adaptation to new environments in asexual organisms.

6 Conclusions

We have compared the adaptation ability of haploid and diploid individuals in different environmental conditions and with different mutation rates. Individuals are simulated agents that interact with an external environment through their simulated neural system. We showed that diploid individuals present more variability in fitness among the individual of a population than haploids and are better able to tolerate environmental changes.
Some results of our simulations, despite the enormous simplification of the model with respect to real biological organisms, seems to confirm some biological data. However, a deeper theoretical analysis is necessary to understand the implications, if any, of our results for biology (for an interesting discussion on a potential cross-fertilization between Artificial Life and Evolutionary Biology see Toquenaga and Wade, 1996).
A straightforward extension of our research on which we are working is the introduction of sexuality into our model, that is, to include a sexual reproduction process with gametogenesis in the model we described.

Acknowledgments

The authors would like to thank the Research Group on Artificial Life of the Italian National Research Council, Valerio Sbordoni of the University "Tor Vergata" and Günter P. Wagner of the Yale University for useful discussions and suggestions.

References

Buss, L. W. (1987). *The evolution of individuality*. Princeton University Press, Princeton, NJ.

Calabretta, R., Galbiati, R., Nolfi, S. & Parisi, D. (1996). Two is Better than One: a Diploid Genotype for Neural Networks. Neural Processing Letters, **4**, 1-7.

Cangelosi, A., Parisi, D. & Nolfi, S. (1994). Cell division and migration in a 'genotype' for neural networks. Network, **5**, 497-515.

Collingwood, E., Corne, D. & Ross, P. (1996). Useful Diversity via Multiploidy. In *Proceedings of IEEE 3rd International Conference on Evolutionary Computation*, Nagoya, Japan.

pard Crow, J. & Kimura, M. (1965). Evolution in Sexual and Asexual Populations. *Am. Nat*., **XCIX** (909), 439-450.

Dellaert, F. & Beer, R. D. (1994). Toward an Evolvable Model of Development for Autonomous Agent Synthesis. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the fourth International Conference on Artificial Life*. MIT Press, Cambridge, MA.

Fonteix, C., Bicking, F., Perrin, E. & Marc, I (1995). Haploid and Diploid Algorithms, a New Approach for Global Optimization: Compared Performances. *Int. J. Systems Sci*., **26**, 1919-1933.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, MA.ar

Goldstein, D. B. (1992). Heterozygote Advantage and the Evolution of a Dominant Diploid Phase. *Genetics*, **132**, 1195-1198.

Holland, J. J. (1975). *Adaptation in Natural and Artificial Systems.* University of Michigan Press, Ann Arbor, MI.

Kondrashov, A. S. & Crow, J. F. (1994). Haploidy or Diploidy: Which is Better? *Nature* , **351**, 314-315.

Menczer, F. & Parisi, D. (1990). A Model for the Emergence of Sex in Evolving Networks: Adaptive Advantage or Random Drift? In F. Varela & P. Maes (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. MIT Press, Cambridge, MA.

Mitchell, M. & Forest, S. (1994). Genetic Algorithms and artificial life. Artificial Life, **1**, 267-290.

Mondada, F., Franzi, E. & Ienne, P. (1993). Mobile Robot Miniaturisation: a Tool for investigation in control algorithms. In: *Proceedings of the third International Symposium on Experimental Robotics, Kyoto, Japan.*

Nolfi, S., Floreano, D., Miglino, 0. & Mondada, F. (1994). How to Evolve Autonomous Robots: Different Approaches in evolutionary robotics. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 190-197). MIT Press, Cambridge, MA.

Nolfi, S. & Parisi, D. (1995). 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, MA.

Orr, H. A. & Otto, S. P. (1994). Does Diploidy Increase the Rate of Adaptation? *Genetics*, **136,** 1475-1480.

Otto, S. P. & Goldstein, D. B. (1992). Recombination and the Evolution of Diploidy. *Genetics*, **131**, 745-751.

Paquin, C. & Adams, J. (1983). Frequency of Adaptive Mutations is Higher in Evolving Diploid than Haploid Yeast Population. *Nature*, **230**, 495-500.

Perrot, V., Richerd, S. & Valèro, M. (1994). Transition from Haploidy to Diploidy. *Nature*, **351**, 315-317.

Smith, R. E. & Goldberg, D. E. (1992). Diploidy and Dominance in Artificial Genetic Search. *Complex Systems*, **6,** 251-285.

Strickberger, M. W. (1976). *Genetics*. Macmillan Publishing, NY.

Syslo, M. M., Deo, N. & Kowalik, J. S. (1983). *Discrete Optimization Algorithms with Pascal Programs*. Prentice-Hall, Englewood Cliffs, NJ.

Toquenaga, Y. & Wade, M.J. (1996). Sewall Wright meets Artificial Life: the Origin and Maintenance of Evolutionary Novelty. *Trends in Ecology & Evolution*, **11**, 478-482.

pard Wagner, G. P. and Altenberg, L (1996). Complex Adaptations and the Evolution of Evolvability, *Evolution*, **50**, 967-976.

Wallace, B. (1981). *Basic Population Genetics*. Columbia University Press, New York.