1994-11

Cell division and migration in a 'genotype' for neural networks

Cangelosi, A

<http://hdl.handle.net/10026.1/3632>

10.1088/0954-898X/5/4/005

Network: Computation in Neural Systems

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Cell division and migration in a 'genotype' for neural networks

(Cell division and migration in neural networks)

Angelo Cangelosi Domenico Parisi Stefano Nolfi

Institute of Psychology - CNR
15, viale Marx
00137 - Rome - Italy
+39-6-86090231

email acangelo@crl.ucsd.edu
email domenico@gracco.irmkant.cnr.it
email stefano@kant.irmkant.cnr.it

In press in *Network: computation in neural systems*

Cell division and migration in a 'genotype' for neural networks

ABSTRACT

Much research has been dedicated recently to applying genetic algorithms to populations of neural networks. However, while in real organisms the inherited genotype maps in complex ways into the resulting phenotype, in most of this research the development process that creates the individual phenotype is ignored. In this paper we present a model of neural development which includes cell division and cell migration in addition to axonal growth and branching. This reflects, in a very simplified way, what happens in the ontogeny of real organisms. The development process of our artificial organisms shows successive phases of functional differentiation and specialization. In addition, we find that mutations that affect different phases of development have very different evolutionary consequences. A single change in the early stages of cell division/migration can have huge effects on the phenotype while changes in later stages have usually a less dramatic impact. Sometimes changes that affect the first developmental stages may be retained producing sudden changes in evolutionary history.

1. INTRODUCTION

In recent years much research has been dedicated to applying genetic algorithms to populations of neural networks (Yao, 1993). Individual networks reproduce selectively as a function of their ability to perform some task (fitness) by generating one or more copies of themselves (offspring). Reproduction can be agamic (single parent) or sexual (two parents). Sexual reproduction increases population variability by introducing new individuals which are new combinations of the two parent networks. However, even with agamic reproduction random changes at reproduction (mutations) allow the population to maintain its variability and to explore new solutions.

In most simulations using genetic algorithms with neural networks no distinction is made between the inherited genotype and the resulting phenotype. What is inherited is a literal description of the phenotypic network which is then evaluated for its fitness. In real organisms genotypes are distinct from phenotypes. What is genetically inherited is not the phenotype itself but a set of instructions for constructing a phenotype. Furthermore, the mapping from genotype to phenotype is a complex, nonlinear process (ontogeny or development). The phenotype which will result from a given genotype cannot be read in the genotype or predicted from it (Langton, 1992). This is important because while what is evaluated for its fitness and, therefore, determines the course taken by evolution, is the phenotype, what changes evolutionarily because of mutations and/or sexual reproduction is the genotype. Many significant properties and intricacies of evolution derive from this separation of levels.

Some attempts have been made to define "genotypes" for neural networks and to describe the mapping process which, given a genotype, constructs a phenotypic network. For example, Harp, Samad, and Guha (1989) define an encoding scheme for genotypes which does not specify all the details of a neural network. A genotype is constituted by one or more segments

that contain two sets of instructions: one set specifies the parameters of an "area" (set of units) and the other set specifies where the connections departing from the area will project. Their results suggest that this more complex representation scheme for neural networks may favour the evolution of more efficient neural networks than those typically used in the literature.

Nolfi and Parisi (1992; in press) have proposed a model of the genotype-to-phenotype mapping in which the genotype specifies the location in space of each of a set of neurons and the growth parameters of each neuron's axon. Connections are established when the branching axon of a neuron reaches another neuron. (We will return to this model in the next Section).

Some researchers have made attempts at simulating the cell division component of biological development. The genotype is imagined to be a collection of rules governing the critical mitotic divisions causing a cell to divide into two 'daughter' cells. These genetic transitions can be modelled as grammatical production (a.k.a. rewrite) rules that are recursive in the sense that they can be applied to their own results. (Some pioneer work with this approach was made by plant morphologist Aristid Lindenmayer, 1976. Cf. also Wilson, 1987; Kitano, 1990.)

Gruau (1993; 1994) has proposed the 'Cellular Encoding' method for developing boolean neural networks. In his model the cells divide or are subject to parameter change depending on the grammatical rules. The use of a genetic algorithm to recombine the grammatical rules favours the emergence of modular architectures.

Belew (1993) also describes a grammar-based model for generating the units of a neural network. The genome contains a set of developmental instructions, some of which become "relevant" to the particular context in which each developing cell finds itself. The initial germ cell represents the start symbol and a series of production rules (the genome) specify the expansion (mitosis) of this nonterminal cell into two other symbols. All cells develop according to the same set of genetically-determined rules until all cells are in a mature, terminal state.

The consequences of adopting a recursive mapping like the one just illustrated are several. First of all, we should consider that a recursive mapping can cause the development of very complicated structures on the basis of very simple genetic instructions. (Cf. the simple simulation experiments by Dawkins, 1986). This is a very important property that may allow the search space to be explored by the evolutionary process without exponential increase with increasing complexity of the phenotype. Kitano (1990) has criticized literal encodings for neural networks such as that proposed by Miller, Todd, and Hedge (1989) because they do not scale well and has described an encoding scheme that tries to avoid this problem.

One should also notice that, in a recursive mapping, although the same set of genetic instructions contained in the original gametic cell may control transitions of all its progeny, the differing contexts of the various cells are likely to cause different genetic instructions to be applied to different cells. The result of this is that there will be some genetic instructions that will determine the first steps and others the final steps of the developing process. Changes (due to mutations or other genetic operators) to different genetic instructions would then cause very different effects. Some changes will affect the fundamental initial steps of the development process and therefore will have very drastic consequences while other changes will affect later developmental stages and will have less dramatic consequences for the resulting phenotype.

For example, if a particular non-terminal cell is responsible for the generation, through successive mitosis and differentiation processes, of an entire functional group of terminal

cells, an individual with an additional group of cells of the same type may easily emerge in successive generations through small changes in the genotype. The resulting individual with the duplicated functional group of cells may win the competition with its conspecifics if the additional group of cells happens to have an adaptive function or if it may be successively modified in order to assume such an adaptive function (pre-adaptation). This is an example of a drastic change (i.e. a change that affects the early stages of development) which may result in a viable individual and in addition may allow the sudden emergence of a new structured character that apparently seem to require the accumulation of a huge set of evolutionary changes. Belew, in the model referred to above, used a duplication genetic operator and found such an operator useful in creating a constructive form of redundancy. As Dawkins has noted (Dawkins, 1986), different species of snakes have different numbers of vertebrae, which means that the number of vertebrae must have changed in evolution quite often. Now, to change the number of vertebrae in an animal is much more complex than just adding an extra bone. Each vertebra has, associated with it, a set of nerves, muscles, blood vessels, etc. On the other hand, the middle part of the body of a snake is composed of a number of segments, many of which are exactly like any other. Since there exists genetic machinery for making one snake segment, new identical segments may easily be added by a single or a few mutations.

2. CELL DIVISION AND CELL MIGRATION

In the developmental model described in Nolfi and Parisi (1992; in press) what is simulated is the process of neural growth. Neural networks are viewed as physical objects in bidimensional space and not, as is usually done, as purely topological structures. Neurons are assigned physical positions in the bidimensional space of the nervous system and they grow branching axons. If one of these branches happens to end up where another neuron is located, a connection between the two neurons is established.

In real organisms, however, the growth of neurites (axons, branches, dendrites) is only one aspect of neural development. Neural development does not start with a set of neurons already positioned in their final spatial location, which then grow their axons and establish connections. Neural growth is preceded by other stages of development in which neurons are created by cell division and they move in space to reach their final location. More generally, development starts from a single cell and it goes through a stage of cell division and cell migration. With regard to nervous cells, neurite growth usually follows completion of the cell division stage (mitotic arrest) and the reaching by individual cells of their appropriate location in the nervous system's 3D-space.

In this paper we present some simulations which add a cell division and cell migration component to the neural growth process described in Nolfi and Parisi (1992). The simulations use an ecological approach to neural networks (Parisi, Cecconi, and Nolfi, 1990). Networks are viewed as models of the nervous system of simple organisms that live and evolve in an environment. It is the environment and the behavior of the organism in the environment which at any given time determine which input is received by the network. We have selected an environment and an evolutionary task (the task in terms of which the fitness of each individual is assessed) that allow us to better explore the potentialities of our model of neural development. The environment contains food and water. The task is to eat when the organism is hungry and to drink when it is thirsty. Such a task may require a modularized network architecture with separate modules (subnetworks) for processing information about food location and about water location. We want to examine how such complex network architecture emerges evolutionarily given our model of neural development which includes a cell division and migration component and a neurite growth component.

3. DESCRIPTION OF THE MODEL

Each individual starts as a single cell. The egg cell contains three kinds of information:

(1) The type of the cell. There are 16 cell types.

(2) A number of parameters concerning the growth process of a neuron:

- the angle of branching of the neuron's axon (there are five branching cycles; each branching is binary; all branchings have the same angle);
- the length of the branching segment of the neuron's axon;
- the "face" of the neuron, i.e. the point on the neuron's surface where the axon grows;
- the bias (i.e. threshold) of the neuron;
- the weight of the connections departing from the neuron (all the connections departing from the same neuron have the same weight).

(3) A set of 16 rules of cell reproduction. Each rule has the following form:

Rewrite Type N as Type N' + Type N".

N goes from 1 to 16. An example of a rule is the following:

Rewrite Type 5 as Type 2 + Type 14.

When a cell divides, it applies the rule which corresponds to its type. In our example, a cell of Type 5 is replaced by two daughter cells which are of Type 2 and Type 14, respectively.

In addition to specifying the type of the two daughter cells, a reproduction rule also specifies two additional informations (cf. Figure 1):

- a number of changes to be made to the neuron parameters (cf. (2) above) when they are inherited by each of the two daughter cells (the changes need not be identical for the two daughter cells);
- the location of the two daughter cells in the bidimensional space of the nervous system. The location of a daughter cell is specified relative to the location of its mother cell. More concretely, for each daughter cell it is specified which of the 8 locations around the mother cell's location will be occupied by the daughter cell. (Neural space is a grid of little squares. Each cell occupies a little square in the space.)

insert figure 1

The process of development begins with the egg cell being located approximately in the center of the space. There are five cell division cycles. Therefore, at the end of the mitotic (division) phase the organism's nervous system will include 32 cells (neurons). (The neurons can be less than 32 if some reproductions rules rewrite a mother cell into a single daughter cell or even no cell at all.)

The 32 neurons end up in spatial positions which depend on the type of the initial egg cell and on the reproduction rules, more specifically on the information contained in each rule specifying the location of daughter cells with respect to the location of mother cells. Neural space is divided into three horizontal bands. A neuron which ends up in the lower band will function as one of the network's input (sensory) units (or motivational units; see below). A neuron which ends up in the upper band will be an output (motor) unit. All the neurons which end up in the intermediate band are hidden units. What an input or output neuron actually codifies is specified in the genetic information associated to each cell type.

At this point the neurite growth process can begin. Each neuron will grow its branching axon according to the values of branching angle and segment length as specified within the neuron itself (see above). There are 5 branching cycles for all neurons. If a branching axon reaches another neuron, the two neurons become connected. Otherwise, branchings will remain as nonfunctional parts of the phenotypic network. Also, subnetworks connecting two or more neurons but which do not map the network's sensory input into the network's motor output, are nonfunctional. All the connections departing from a neuron are assigned the same weight and each neuron has its own bias. The resulting functional network controls the organism's behavior in the environment during the organism's life.

An initial population of 100 egg cells are randomly generated. Everything in each egg cell is randomly decided: cell type, the values of the neural growth parameters, the set of 16 rules. Then development begins (cell division, cell migration, neurite growth) and it eventually results in 100 'adult' neural networks. (Development is instantaneous. An organism begins to behave in the environment as an adult.) The 100 networks live for a fixed number of cycles of activation spreadings in their environment. At the end of life they reproduce selectively as a function of fitness (see next Section).

Each reproducing individual generates a number of copies of its egg cell, one for each offspring. Random mutations are applied to all the information contained in the egg. Hence, an offspring network will be similar but not identical to its parent network and to its sibling networks.

4. THE ENVIRONMENT AND THE TASK

The environment and the task used in the present simulations are those of a previous simulation described in Cecconi and Parisi (1993) with some modifications. Furthermore, although a genetic algorithm was also applied in the Cecconi and Parisi's simulation, in that simulation network architecture was fixed and did not evolve. Only weight matrices varied from one individual to another and were genetically inherited. The genotype/phenotype mapping was one-to-one. The phenotypic network was directly described in all details in the genotype.

Each organism lives alone in an environment which is a bidimensional space of 40x40 little squares. The environment includes two separate randomly placed zones (2x2) one containing food and the other water. The organism occupies a single square and it has a facing direction. It can move one cell ahead but if it wants to go in another direction it must first turn left or right 90 degrees. At each cycle the organism is informed by its senses of the location of the food zone and of the water zone. The organism's task is to reach the food zone and to remain there when it is hungry and to reach the water zone and remain there when it is thirsty. In other words, the organism's fitness is increased when food is eaten but only if the organism eats while being in a state of hunger - and the same for water and thirst.

The organism's behavior is controlled by a neural network which is the terminal result of the organism's neural development as described in the preceding Section. The network may include input units, output units, hidden units, and motivational units. The input units encode the angle and the distance of the (center of) the food zone and the angle and distance of the (center of) the water zone. Angles are measured clockwise with respect to the organism's current facing direction and mapped in the interval between 0 and 1. The distance is the Euclidean distance from the organism, also mapped in the interval between 0 and 1. Two different types of output units binarily encode four possible motor actions: 11=go one step forward in the facing direction; 10=turn 90 degree on the left; 01=turn 90 degree on the right; 00=do nothing. Motivational units encode the motivational state of the organism (see below). All the remaining units are hidden units that map the activation state of the input and motivational units into the activation state of the output units.

The motivational units encode a state of hunger or a state of thirst in a binary fashion. The organism remains in a certain motivational state until it has spent 10 time steps in the food or water areas. After that, the motivational state abruptly changes to the opposite state. Each motivational units encodes the motivational state (hunger = 1; thirst = 0) independently from the other motivational units. If there are more than a single motivational unit, they each independently influence the network's behavior through their connections to the hidden units.

The motivational units function as internal input units. (This is why they are placed in the lower band of the nervous system.) As input units they send their excitations or inhibitions to the hidden units depending on the weights of their connections to each hidden unit. However, unlike the other input units, motivational units are not activated by the external environment but by the internal environment of the organism's body. What is expected to evolve in this population of networks is an ability to use information about the location of the food zone to move in the environment in order to reach the food zone when the motivational state is hunger, ignoring the equally present information about the location of the water zone. And to use information about water to reach water when the motivation is thirst, ignoring food. The fitness of each individual at the end of life is computed by adding the number of time steps spent in the food zone when hungry and the number of time steps spent in the water zone when thirsty, minus the difference between the two. This third component of the fitness formula was introduced to induce a behavior balanced between searching for food and searching for water (cf. Cecconi and Parisi, 1993).

Each individual network lives for a total of 1500 spreadings of activation. Total life is composed of 10 'epochs'. Each epoch is 150 spreadings in a different environment, where an environment is defined by the particular location of the food and water zones. At the beginning of each epoch the organism is placed in a randomly selected position of the environment (including positions inside the food and water zones).

At the end of life some but not all individuals reproduce. The 20 networks with the highest fitness generate 5 offspring each (copies of their egg cell with random mutations) whereas the other 80 networks die without leaving offspring. The 100 new networks constitute the next generation. The process continues for a total of 500 generations.

The fitness of an individual organism is as a property of that organism. Hence, the evolutionary curve of fitness tells us something about how these organisms change during 500 generations. But in our simulations organisms can be analyzed at different levels and changes at one level can be compared with changes at other levels. More specifically, we can analyze our organisms at the genetic, neural, behavioral, and fitness level (cf. Miglino, Nolfi, and Parisi, in press).

5. RESULTS

We have run two sets of simulations. In one set the information contained in the genotype is exactly the information shown in Figure 2 above. In the other set some of the information encoded in the genotype has been eliminated, and precisely the information regarding angle of branching and segment length of neurons' axons and the "face" of neurons. These parameters are fixed and they do not evolve. In this second set of simulations, therefore, network architecture depends entirely on the cell division and migration component of development. Since the results of the two sets of simulations are very similar, we will report the results of the second set only.

insert figure 2

Figure 2 shows the average fitness and the fitness of the best individual for each one of the 500 generations. Two aspects of these curves are to be noted. First, fitness increases evolutionarily but the increase has a typical pattern of sudden jumps followed by long flat periods. The main jumps occur at around Generation 100 and at around Generation 300. (The initial jump at around generation 20 is less important because it is due to the elimination from the population of individuals which simply do not move.) Second, while the pattern of growth is similar for both the average and the best individual curve, there is an increasing distance between the two curves across evolutionary time. We will return to these two results in Section 6.

insert figure 3

Figure 3 presents the evolutionary curve of fitness of the best individuals (same as that of Figure 2) alongside with evolutionary changes in a number of other aspects of our organisms: (a) changes in behavior; (b) changes in the motivational component (motivational units) of neural architecture; (c) changes in network modularization with respect to food and water; (d) genetic changes; more particularly, changes due to mutations of cell division rules.

When we compare changes at these different levels, some of the results that emerge are the following. Total evolutionary time can be divided up into three stages.

Stage 1

The initial flat period during which fitness remains the same until Generation 100 corresponds to a behavior which allows the organism to eat but not to drink. In other words, the best individuals of the first 100 generations are able to reach the food zone but they completely ignore the water zone.

During this period their functional neural architecture does not include motivational units. A food pathway (i.e. a neural module that maps sensory information about food zone into motor actions; cf. below) has already formed at this stage but there is no corresponding pathway for water. Finally, at the genetic level a mutation of the rule for the second cell division appears at the end of the period. This mutation will have many important phenotypic effects in the next stage.

Stage 2

From Generation 100 to Generation 300 a new stage is reached at all levels. Fitness level is higher than in the previous stage. At the behavioral level the organisms continue to be able to reach the food zone but they begin to pay attention to water although their behavior with respect to the water zone still leaves much to be desired. At the neural level a single motivational unit first appears and it is then integrated with the rest of network architecture. At the same time a water neural pathway begins to be formed. At the end of the stage a mutation of the first cell division rule appears in a particular individual which will cause the transition to the third stage.

Stage 3

After Generation 300 the third and final stage is reached. The organisms are now equally able to reach the food zone when hungry and the water zone when thirsty. The neural pathway for processing sensory information from the water zone is completed. Hence, neural architecture tends to include two modules, one for food and one for water.

To better understand the nature of these modules we will examine the network architecture of the best individual of one of the last generations (Generation 440). This architecture is shown in Figure 4 together with its connection weights (on connections) and biases (inside units).

insert figure 4

Various things can be noted about this architecture. First, the input units encode the angle of the food zone and the angle of the water zone but not their distances. Network architectures which encoded both angle and distance or distance only have been eliminated from the population in the preceding generations. This result shows that sensory and motor systems tend to be co-adapted and to co-evolve. Given the primitive nature of the organism's motor behavior (single step forward and 90 degree turns), angle information is sufficient to take the appropriate motor decisions. Distance information is useless (or, worse, is interfering) and is eliminated.

Second, the architecture includes two distinct pathways (modules), one for food and the other for water. These pathways are shown in Figure 5.

insert figure 5

The two pathways have been identified using the following procedure. We have tested the network with all possible stimuli in two conditions: with the motivational unit set to 1 (hunger) and with the motivational unit set to 0 (thirst). For each stimulus (the pair of angles of the organism with respect to the food zone and to the water zone) we have measured the activation level of each of the hidden units and of the average output. What is obtained is the "receptive field" of these units, i.e. how each unit responds (in terms of its activation level) to all possible input stimuli.

insert figure 6

The results are presented in Figure 6. On the left we show the receptive fields of the units when the motivation is hunger; on the right the receptive fields when the motivation is thirst. The receptive field of a unit is represented as the variation of activation level of the unit in a bidimensional space. The x axis of the space represents unit variation associated with variation of the input angle from the food zone (food angle); the y axis represents unit variation associated with variation of the water angle (cf. Legend).

When the motivation is hunger (left side of Figure 6) three units respond to variation in food angle while all the other units are not responsive to input at all. Therefore, when the motivation is hunger (i.e. when the motivational unit has a value of 1) information about water is simply ignored by the network. Only information about food is responsible for what the network does. The three units which respond to information about food when the motivation is hunger constitute the food neural pathway.

When the motivation is water (right side of Figure 6), the units which were indifferent to input variation in the preceding condition, now become responsive to variation of water angle. But the situation is not symmetric with respect to the previous case. Two of the three units which constitute the food pathway appear to have also a role in processing information about water. (Cf. Figure 5b) Hence, the water pathway includes some units which are specialized for processing information about water and some units which are also used to process information about food. This asymmetry appears to be related to the history of the evolution of the two abilities of finding food and finding water. As will be remembered, the ability to find food emerges evolutionarily in this population before the ability to find water.

We can now reconstruct how the architecture of Figure 4 develops in the individual. In particular we will describe the cell division and migration stage of development.

insert figure 7

The five cell division phases are shown in Figure 7. The first picture (7a) shows the egg positioned approximately in the center of the future neural space (i.e. the space that will contain the future nervous system). The space is divided into three horizontal bands. The upper band will contain motor output units. The lowest band will include input units, both sensory and motivational. The larger intermediate band will contain the hidden units.

Figure 7b shows the result of the first cell division (2 cells). The egg has generated two daughter cells, labeled A and B. The nature of these cells and their position in space are dictated by the genetic information contained in the egg.

The next picture shows the result of the second cell division (4 cells). Each daughter cell has generated two new cells, A1 and A2, and B1 and B2, respectively. We can begin to recognize the future role of these cells. The A1 and B1 cells are founder cells for hidden and output units. This means that their descendants will mostly be either hidden or output units. The B2 cell is a founder cell for input units. Its descendant will mostly be input units. The A2 cell, on the other hand, is a founder cell for nonfunctional units. Its descendants will be part of the nonfunctional network.

The results of the third cell division are shown in the next picture (8 cells). Cells continue to differentiate and specialize. The labels in the picture indicate these specializations. For example, the B2 cell, which was a founder cell for input units (both sensory and motivational input), divides up into two cells, B21 and B22. From B21 the input units encoding food and water position will be derived, while B22 will originate the motivational input unit.

The fourth (16 cells) and fifth (32 cells) divisions are shown in the next two pictures. No labels are included in these two pictures for practical reasons. The last two pictures show the results of the axonal growth phase after the last cell division (7g) and the resulting functional network after the growth of branching axons and the establishment of connections (fh). This functional network, of course, is the same network of Figure 4 (with some changes in spatial relations for purposes of presentation).

The entire cell lineage is summarized in Figure 8. This figure shows cells labels (same labels used in Figure 7), cell types, the function of the terminal cells (input, output, motivational, and hidden units), and the two network modules (the food and water pathways).

insert figure 8

6. DISCUSSION AND CONCLUSIONS

We can return to the two aspects of the evolutionary curves of fitness increase that we have already noted. First, the curve shows a pattern of long flat periods with no fitness increases punctuated by sudden large increases in fitness and, second, there is a large and increasing distance between the average and the best fitness value within each generations.

How to explain punctuated equilibria is highly controversial. Gould (1991) claims that the idea that all change is gradual and continuous is a consequence of interpreting changes as mostly adaptations while most of retained changes are neutral. The accumulation of neutral changes may produce pre-adaptation phenomena because retained neutral characters may later turn out to be adapted to new circumstances. Pre-adaptations may in turn explain punctuated equilibria, e.g. because when a neutral character becomes a useful ones through a new additional change, the fitness and/or behaviour can suddenly change.

Miglino, Nolfi, and Parisi (in press) show that models of development that assume a complex genotype/phenotype mapping and, therefore, a complex hierarchical organization of levels within organisms, tend to yield fitness curves with a punctuated equilibrium shape. Many changes are likely to accumulate at lower levels without being expressed at higher levels. This results in stasis. A further small change at a low level (e.g. a single mutation at the genetic level) may interact with already accumulated changes to determine large changes at higher levels. This results in sudden jumps.

It is clear that as models of developmental mapping become more complex, these phenomena of hidden changes at lower levels that become suddenly visible at higher levels will tend to increase in importance. In the present model a cell division and migration stage is added to the axonal growth stage already present in the model of Miglino, Nolfi, and Parisi (in press). A single change in the first phases of cell division/migration can have huge consequences on the phenotype. If hidden accumulated changes concern the first phases of development, when they become visible at higher levels because of a further small change they can have large

phenotypic consequences at the neural, behavioral, and fitness level. This is what is observed in the results of our simulations.

We wish now to comment on the increasing distance between average and best fitness that we find in the present simulations. This increasing distance implies that the distribution of fitness values within a generation tends to have an asymmetrical shape. It is a right skewed distribution with many low values and a long tail of high values. In other words, evolution succeeds in finding some very good solutions to the task (the few high fitness values) but most solutions are not very good. This picture is different from the picture that usually emerges in simulations of the evolution of neural networks, where the distribution of fitness values within generations tends to be more nearly symmetrical (Gaussian) even in later generations and, therefore, the distance between best fitness values and average values is less large (or even there may be an evolutionary tendency towards a reduced distance between average and best fitness).

We believe that the great distance between average and best fitness in the present simulation may be related to the particular model of development used in the simulation. In simulations with the same task but a simpler one-to-one model of development (i.e. phenotypic networks are literally encoded in the genotype) the distance between average and best fitness in later generations tends to be much smaller than in the present simulation (Cecconi and Parisi, 1993). This is shown in Figure 9 which compares the distribution of fitness values within a single generation (100 individuals) in the simulation with a simple one-to-one mapping and in the current simulation with a more complex mapping. The two distributions have very different shapes.

insert figure 9

A possible explanation of this difference might be the following. Typical simulations of the evolution of populations of networks that use mutations as the sole genetic operator, tend to use high mutation rates (e.g. 10% of the connection weights are mutated). These high rates seem to be necessary to create enough variability for selection to be successful (especially if sexual recombination is not used, as in the present case). However, mutation rate may interact with the particular model of development which is used. More specifically, the effect of a single mutation on the phenotype (the network) can be much greater if the mapping from genotype to phenotype is complex and many-to-many than if it is a simple one-to-one mapping.

Consider the present model. As already noted, a single mutation that happens to change the genotype at the level of the cell reduplication rules can have a huge effect on the resulting network. For example, if a mutation changes the cell type of one of the two daughter cells in the first cell division during the development of an individual, such mutation can change a whole half of the offspring's network with respect to its parent's network. This is not the case if the genotype/phenotype mapping is one-to-one. A single mutation can only have an effect on a single feature of the phenotypic network; for example, it could eliminate or create a connection or it could change a single connection weight. The developmental model used in Miglino, Nolfi, and Parisi (in press), which is more complex than the one-to-one mapping model used in most current simulations but does not include cell division and migration, is probably intermediate between the two. A single mutation in that model can have more phenotypic consequences than in a one-to-one mapping model but not as many as in a model with cell division and migration.

In the present simulation with a complex developmental mapping which includes cell division and migration we have used more or less the same mutation rate as in our previous simulations (10%). But the phenotypic effects of a single mutation can be much greater in the present case. This might explain why later generations include a few very good individuals and a large number of not very good individuals. Large effects of single mutations can be disruptive. While some individuals succeed in inheriting an increasingly better phenotypic network, the majority of individuals within a generation have their inherited network disrupted by the huge effects of mutations. Hence, the distance between average and best fitness values tends to be great in later generations.

The morale of this history is that a possible criticism of current simulations on evolving populations of neural networks, i.e. that they use unrealistically high mutation rates compared with the mutation rates found in nature, can be answered if one takes the type of developmental model used in the simulations into consideration. Developmental mappings in real organisms are likely to be closer to our complex developmental model that includes cell division and migration than to simple one-to-one mapping models. But single mutations can have much larger phenotypic effects in complex mappings than in simpler mappings. This might explain why mutations rates can be so low in nature and at the same time natural selection can have enough variability to work with. The high mutation rates and variability found in such simple organisms as viruses (Eigen, 1993) appear to be compatible with this idea since viruses have an extremely simple genotype-to-phenotype mapping.

In conclusion, the model of neural development described in this paper, which includes cells division and cell migration in addition to axonal growth and branching, is organized in successive phases of functional differentiation and specialization of the artificial neurons. This reflects, in a very simplified way, what happens in the ontogeny of real organisms. The development of neural structures begins with a folding in of the ectodermic tissue which forms the neural crest (Purves and Lichtheim, 1985). This structure gives origin, on the basis of a progressive differentiation into cytologically and functionally different regions, to various structures of the nervous system. The differentiation process has three phases: proliferation and genesis of different classes of neurons, migration, and cell maturation with stabilizations of synapses (Schacher, 1985). These basic properties of neural development appear to be all present in the ontogenetic development of our neural networks.

The model we have described is extremely simplified with respect to real biological organisms. Neurons are placed in a two-dimensional space instead of the three-dimensional space of real neural systems. Moreover, since the total number of allowed mitoses is 5, a network can be composed of a maximum numbers of 32 neurons. Despite of that, the model allows us to study aspects of evolutionary change at the genetic, neural, and behavioral level and, especially, their interactions. Because of the complexity of each level traditional disciplines tend to concentrate on single levels and in many cases miss these interactions. The method of computer simulation necessarily leaves out many possibly important details but it offers the advantage of examining how change at one level can be related to change at other levels.

References

- Belew R K 1993 Interposing an ontogenetic model between genetic algorithms and neural networks *Advances in Neural Information Processing (NIPS5)* ed J Cowan (San Mateo, CA: Morgan Kaufmann)
- Cecconi F and Parisi D 1993 Neural networks with motivational units *From Animals to Animats 2: Proceedings Second Intl. Conf. on Simulation of Adaptive Behavior* eds J Meyer H L Roitblat and S W Wilson (Cambridge, MA: MIT Press, Bradford Books)
- Dawkins R 1986 *The Blind Watchmaker* (New York: Norton & Cooper)
- Eigen M 1993 Viral quasispecies *Scientific American* **269** 32-39
- Gould S J 1991 Exaptation: a crucial tool for an evolutionary psychology *Journal of Social Issues* **3** 43-65
- Gruau F 1993 Genetic synthesis of modular neural networks *Proceedings Fifth Intl. Conf. on Genetic Algorithms* (San Mateo, CA: Morgan Kaufmann)
- Gruau F 1994 Genetic Micro Programming of Neural Networks *Advances in Genetic Programming* ed K E Kinnear (Cambridge, MA: MIT Press, Bradford Books)
- Harp S, Samad T and Guha A 1989 Toward the genetic synthesis of neural networks *Proceedings Third Intl. Conf. on Genetic Algorithms* (San Mateo, CA: Morgan Kaufmann)
- Kitano H 1990. Designing neural networks using genetic algorithms with graph generation system *Complex Systems* **4** 461-476.
- Langton C G 1992 Artificial Life *Lectures in Complex Systems* eds L Nadel and D Stein (Reading, MA: Addison Wesley)
- Lindenmayer A and Rozenberg G 1976 *Automata, Languages, Development* (North-Holland, Amsterdam)
- Miglino O, Nolfi S and Parisi D in press Discontinuity in evolution: how different levels of organization imply pre-adaptation *Plastic Individuals in Evolving Populations* eds R K Belew and M Mitchell (Reading, MA: Addison-Wesley, SFI Series)
- Miller G F, Todd P M and Hedge S U 1989 Designing neural networks using genetic algorithms *Proceedings Third Intl. Conf. on Genetic Algorithms* (San Mateo, CA: Morgan Kaufmann)
- Nolfi S and Parisi D 1992 Growing Neural Networks *Technical Report* (Institute of Psychology, CNR Rome)
- Nolfi S and Parisi D in press Genotypes for neural networks *The Handbook of Brain Theory and Neural Networks* ed M A Arbib (Cambridge, MA: MIT Press, Bradford Books)
- Parisi D, Cecconi F and Nolfi S 1990 Econets: neural networks that learn in an environment *Network* **1** 149-168.
- Purves D and Lichtheim J W 1985 *Principles of Neural Development* (Sunderland, MA: Sinauer Ass.)
- Schacher S 1985 Determination and differentiation in neural system development *Principles of Neurosciences* eds E R Kandel and J H Schwartz (New York: Esvier Science)
- Wilson S W 1987 The genetic algorithm and biological development *Proceedings Second Intl. Conf. on Genetic Algorithms*. (Hillsdale, NJ: Erlbaum)
- Yao X 1993 Evolutionary artificial neural networks *International Journal of Neural Systems* **4** 203-222

List of Captions of Illustrations

Figure 1 - Rules contained in the genotype

Figure 2 - Evolutionary change in fitness during 500 generations. Average fitness and the fitness of the best individual are shown.

Figure 3 - General view of results according to different levels of analysis

Figure 4 - Functional architecture of the best individual of Generation 440. Connections weights and units biases are indicated.

Figure 5 - Modular organization of network. (5a) Food pathway; (5b) Water pathway

Figure 6 - Receptive fields of single network units. (6a) Motivation = hunger; (6b) Motivation = thirst.

Figure 7 - Five cycles of cell division and axonal growth phase. (7a - Egg position. 7b - First mitotic division. 7c - Second mitotic division. 7d - Third mitotic division. 7e - Fourth mitotic division. 7f - Fifth mitotic division. 7g - axonal growing. 7h - functionally connected cells)

Figure 8 - Cell lineage

Figure 9 - Distribution of fitness values in two simulations. (9a) Simulation with fixed architecture neural networks; (9b) Simulation with developing neural networks.

Illustrations

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5a

Figure 5b

Figure 6

7a

7b

7c

7d

7e

7f

7g

7h

Figure 7

Figure 8

Figure 9a

Figure 9b