

MIND AS EVOLUTION

AND EVOLUTION AS SUCH

Robert Pallbo

*Lund University Cognitive Science,
Kungshuset, Lundagård,
S-222 22 Lund,
Sweden.*

E-mail: robert.pallbo@fil.lu.se

Abstract. One of the purposes of consciousness is to acquire knowledge. The models of this task traditionally assume a division between the known and the knower, a division that leads into doubtful reasoning about some kind of (frequently disguised) homunculus. The current paper tries to avoid such fallacies by viewing the adaptation of the mind as an evolutionary process, a process distinct (but not independent of), and faster than phylogenetic evolution. To allow the modelling of the mind as evolution, the first aim of the paper is to outline a general framework of evolution. Within this framework, phylogenetic and ontogenetic (individual) evolution can be generated as special cases. The particular characteristics and mechanisms that allows an evolution of the individual mind is then presented. In this, spontaneous, or stochastic, activity among the neurons take an important position as the variation mechanism. Furthermore, ‘resonance’ between the sensorial input and the brain is taken as the selection mechanism. Finally, this view of the mind allows a simpler explanation of how the mind could evolve in nature since the model does not assume that the mind inherits a huge amount of specialized computational abilities.

1 INTRODUCTION

In the development of the species, we may distinguish two types of adaptation; phylogenetic, and ontogenetic. Phylogenetic adaptation is the way in which a *population* changes its traits in response to an environment. The change is not caused by an adaptation of the individual organisms, but by altering the composition of the group. Ontogenetic adaptation, in turn, is the way in which an *individual organism* adapts to its personal environment. Such adaptation causes the organism to change its behaviour.

A question that emerges in the observation of adaptation is *who* or *what* is doing the adapting (Sober,

1993). In phylogenetic adaptation, the common position is that there is no such doer and there is no need for it. The adaptation that takes place is a consequence of natural selection in an evolutionary process.

The case of ontogenetic adaptation appears at a first glance as much simpler. It is very tempting to say that it is the individual organism that is doing the adapting. Such a statement is, however, empty. This becomes clear if we alter the previous question to ‘exactly who or what *in the organism* is doing the adapting?’ Written like this, the answer is much less obvious.

In the dealing with this difficulty, it is easy to resort to the idea that there is something inside the individual that is the doer. But to separate knowledge from something that knows introduces new difficulties. If knowledge and knowing are not the same, then the knower must by some means access the knowledge, that is to say, interpret it. This, in turn, means that the knowledge must be described in a representation that is known *a priori* by the knower, which is to say that we (our minds) are restricted in what we can learn.

Even if we do not feel satisfied with such an approach, it is hard to see how there could be any alternatives. But there are.

Every culture is based on assumptions so taken for granted that they are barely conscious, and it is only when we study highly different cultures and languages that we become aware of them. Standard average European (SAE) languages, for example, have sentences so structured that the verb (event) must be set in motion by the noun (thing) – thereby posing a metaphysical problem as tricky, and probably as meaningless, as that of the relation of mind to body. We cannot talk of “knowing” without assuming that there is some “who” or “what” that knows, not realizing that this is nothing more than

a grammatical convention. The supposition that knowing requires a knower is based on a linguistic and not an existential rule, as becomes obvious when we consider that raining needs no rainer and clouding no cloudier. (Watts, 1975, p. 11)

Just as in phylogenetic adaptation, where we have dismissed the idea of a something that is doing the adaptation, we can dismiss the doer in the case of ontogenetic adaptation. After all, we already have a theory for it – evolution. The problem is, however, that this theory is adapted to the special case of phylogenetic adaptation. Therefore, the first aim of this paper is to outline a framework of evolutionary processes in general. Thereafter, ontogenetic adaptation will be described as a special case within this framework, that is, the goal is not to explain ontogenetic adaptation within phylogenetic evolution, but as an evolutionary process on its own. Finally, this work will be related to other works in the field.

2 A FRAMEWORK OF EVOLUTION

Since the theory of evolution is adapted to phylogenetic adaptation, we could start from there and reduce the theory until we reach the most general case that we can come up with. Thereafter, we could identify the specifics of ontogenetic adaptation and create a new category of evolution. I have not chosen this way, however. Instead I will present a general framework first. From this general framework several hypothetical cases of evolution will be generated. One of these fits into phylogenetic adaptation, another into ontogenetic adaptation, and still others that we cannot fit into any process. I believe that by presenting it this way, the relationship and, especially, the difference between phylogenetic and ontogenetic adaptation will stand out more clearly.

2.1 A GENERAL PICTURE

An evolutionary process must always take place somewhere. Let us call this place the *universe*. In the development of the species, the universe corresponds to the physical reality. In other evolutionary processes, the universe may be something else, like a virtual world simulated in a computer, or a world that emerges inside a complex entity like the brain. The characteristics of the universe will in a wide sense characterize the evolutionary processes that it is hosting.

In the universe, there must be a ‘something’ on which the evolution can act, i.e., some kind of *substance*, which may be concrete, as matter in the case of the species, or abstract, as ideas in the case of cultural evolution. If there is an evolutionary process acting on this substance, then the substances will, as a

result, form complex *structures*. These structures are the things that are adapting. In the evolution of the species, the structure corresponds to the species themselves.

To allow the structure to adapt, we need to introduce two mechanisms; *variation* and *selection*. Variation is a random, or ‘blind,’ mechanism that modifies the structure. By ‘blind,’ I mean that the variation mechanism does not know about the existence of the structure. Rather, the variation mechanism is a property of the substances that make up the structure and, hence, is, so to say, inherited from the universe. That we speak of the variation as ‘acting’ on the structure, is an external abstraction made solely for the purpose of simplicity in discussion.

The second mechanism, selection suggests by its name, that the best parts of the structure are to remain while the rest are to be abolished. Such an operation, however, requires that the mechanism knows what is to be regarded as ‘good,’ and what is not. Obviously, such an idea is absurd. Selection is not an intentional (nor a ‘blind’) *selector*.

What the variation mechanism creates, is not necessarily a stable structure configuration in the environment of the structure (I will discuss more about the environment below). An unstable configuration will vanish or be restored to a stable configuration. When the variation mechanism has caused a stable structure, however, that structure will remain. What is stable and what is not, is determined both by the substance itself that the structure is made of, as well as the specific environment in which it resides. The environment can be said to carve out the structure from substance modifications caused by the variation mechanism. This carving is what we call selection. Note that, like the variation mechanism, the selection mechanism does not need to know about the structure.

To sum up, the variation constructs a structure under the supervision of the selection mechanism. Whenever the variation mechanism causes a modification of the structure that is not recovered by selection, we note a ‘development’ of that structure. When the variation continues from this ‘developed’ structure and causes several successive ‘developments,’ we say that the structure is evolving in an evolutionary process. The situation is illustrated in Figure 1a.

2.1.1 Dividing the Structure

In phylogenetic adaptation, we observe that the structure, i.e., the species, can be divided into two components; the phenotype and the genotype.¹ Below, we will make a similar division of the structure involved in ontogenetic adaptation. In that case, the two components corresponds to the neural activity pattern and the

¹With the terms ‘genotype’ and ‘phenotype’ I refer to a group of organisms sharing a pool of genes.

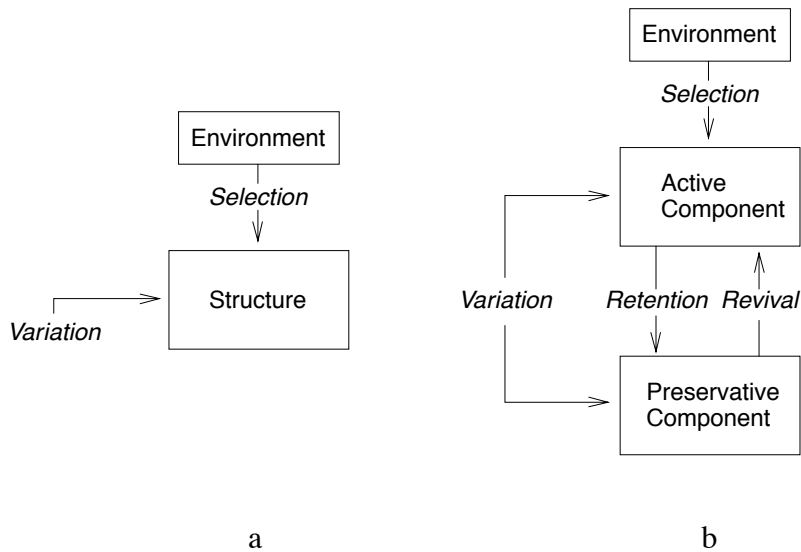


Figure 1: The relations between the components and the mechanisms of an evolutionary process. a) The relation between the structure and the environment. b) The same situation when the structure is divided in two components.

synaptic configuration. What these two examples have in common is that one of the components, the genotype and the synaptic configuration respectively, is exposed to the environment only indirectly through the other component. For this reason, let us call this less exposed component the *preservative* and the other one the *active* component.

The two components are mutually dependent. For this reason, it is useful to introduce two new mechanisms. The first mechanism lets the active component influence the preservative. Every active structure component that is stable in the environment will influence the preservative component. A non-stable active component, however, will vanish before it has had any considerable influence on the preservative component. Consequently, the preservative component comes to accumulate influences from ‘successful’ (or ‘stable’ to use a more neutral word) structures. We will denote this relation as the *retention* mechanism.

The second mechanism is the influence of the preservative component on the active component. The accumulated influences from stable structures into the preservative component are useful only if the preservative component in turn influences the active component back again. If not, the division of the structure into the two components would not contribute to the process and, therefore, probably not take place. The second mechanism, which we will call the *revival* mechanism, must therefore allow a recycling of the accumulated ‘knowledge.’

In phylogenetic adaptation, the retention mechanism involves reproduction of the organisms. Only the organisms that are stable in the environment are able to contribute to the production of new organisms.

Therefore, death and reproduction cause modifications of the current genotype and, hence, there is an influence from the active component (the phenotype) to the preservative (the genotype).

The reproduction causes not only a modification of the genotype, but also triggers the genotype to modify the phenotype. That is, a new organism is created by means of ontogenesis. This part of the picture constitutes the revival mechanism in the evolution of the species.

A final remark needs to be made concerning the division of the structure into two components. It concerns how the variation and the selection mechanisms need to be adapted to this new situation. It has already been said that it is the active component that interacts with the environment. Therefore the selection mechanism will act upon this component as it is this mechanism that constitutes the influence of the environment on the structure.

The variation mechanism, on the other hand, should affect the preservative component as this component serves as a form of ‘memory’ in the process. This is also the case in phylogenetic adaptation where mutation acts directly on the genes to allow variation. But it would also be possible for the variation mechanism to affect the active component as the active component then can affect the preservative component further by means of the retention mechanism. This turns out to be the case in ontogenetic adaptation that will be described below. Finally, it would be possible for the variation mechanism to act upon both the active as well as the preservative component. I have not been able to find any example where this is so, however. The new scenario is summarized in Figure 1b.

2.1.2 *The Universe*

A prerequisite for a structure to be formed is the existence of the substance from which it is composed. This substance, as previously mentioned, exists in some space, which we have denoted as the universe. Furthermore, the substance must possess qualities that can constitute the mechanisms needed in an evolutionary process. These qualities can be seen as a ‘law of Nature’ in the universe.

The substance composing a structure may very well be structures formed by other evolutionary processes. In this way, we can achieve a hierarchy of nested evolution. For instance, the evolution of culture is based upon the evolution of the mind. What is important for this to be possible, is that the structure that constitutes the substance at a new level can inherit, or perhaps form new, qualities that can constitute the evolutionary mechanisms at the new level. When there is no opportunity for this to happen, the hierarchy of evolutionary processes has reached a dead end.

The degree of hierarchical complexity of evolutionary processes depends significantly upon the universe in which they take place. In order to support evolutionary processes, the universe must be rich in terms of substance and opportunities for mechanisms to appear. Furthermore, in any universe that has such a richness, there is no reason to assume that evolutionary processes should *not* emerge.

2.1.3 *The Environment*

The area in the universe that a structure inhabits is usually limited. This limited area, the *environment* of the structure, may exhibit unique conditions. For instance, the conditions above the surface of the sea are much different from those below. An evolving structure will adapt to such local conditions.

The locality in space may not only be understood as a geographical area. We can also localize a structure at a certain place on the *scale* of space and time (cf. Havel, 1993). That is, on the scale of space, humans are located at around one meter. At this scale of space, the world is much different from the same geographical place seen from the perspective of about one micrometer. Each position on the scale has its own conditions, and an evolving structure will adapt not only to the geographical environment, but also to the ‘scale environment’ that it inhabits.

2.1.4 *Saturated Evolution*

An evolutionary process may reach a point where no more adaptations are possible. This may have several reasons. One possibility is that the source of variation is drained. In this case, the evolutionary process will halt. If the structure is stable enough (and it is likely to be so when the source of variation has disappeared) the

structure may be able to maintain itself without evolving. The source of variation may also gradually fade. In this case, the evolutionary process will slow down proportionally.

Another possibility of saturated evolution is that the structure has reached maximum adaptation in the environment. This is unlikely to happen in a complex universe hosting several interacting, and competing, structures, but may very well occur in a simple (read: artificial) universe.

Finally, the structure may have made use of all available substance. Without additional substance, the structure cannot grow (but may, of course, still be modified). This does not exclude the possibility that the saturated structure constitutes the substance to a new level of evolution. In fact, when a new evolutionary process spawns from another, it is beneficial if the underlying evolution has saturated. The reason for this is that the structure from a saturated evolution constitutes a *stable* substance for the new process. If a new evolutionary process spawns from a non-saturated evolution, the new process must somehow stabilize the old process by other means. That is, the underlying process must somehow benefit from the new level (cf. Plotkin and Odling-Smee, 1981).

2.2 DISTRIBUTED EVOLUTION

So far, we have talked about *one* structure and how this structure can be divided into two components. I have avoided dealing with the fact that the species, the structure in phylogenetic adaptation, is composed of individuals. The reason for this inattention is that I believe individuals are optional in an evolutionary process. Nothing, in theory, stops us from assuming that the structure, or the structure components, in the discussion above are solitary entities.

We thus learn from phylogenetic adaptation that the structure, and the two components, may be a compound of individuals. This arrangement will be denoted as *distributed evolution*. We will recognize that there is more than one way in which the structure can be divided into individuals.

When the structure is not divided into an active and one preservative component, the structure is simply *one* set of individuals (Fig. 2a). This is the case of self-replicating molecules.² When the structure *is* divided into the two components, the situation becomes more complex. First, only one of the components may be a compound while the other is solitary. As a second alternative, the active and the preservative component may be distributed together. This is the case in phylogenetic adaptation where each organism constitutes one unit of the active component (the phenotype) and

²Self-replicating molecules were thought to have been discovered (Tjivikua et al., 1990; see also Dagani, 1992). The result was, however, given a simpler explanation without the assumption of self-replication by Menger et al. (1994).

one unit of the preservative component (its personal set of genes). Finally, the active and preservative components may be distributed individually. This is the case in ontogenetic adaptation where the neural activity pattern (the active component) is distributed over the neurons, and the synaptic configuration (the preservative component) is distributed over the synapses. All these varieties of distributed evolution are summarized in Figure 2.

When the active and the preservative structure components are distributed together as is the case with the species (Fig. 2b), it is important that the mechanisms are acting locally upon and within the individuals. The retention mechanism, however, must not be isolated within a single individual and affect its personal structure alone. The retention mechanism must operate on the preservative component set as a whole, or we would not have *one* distributed evolution but many, independent, and non-distributed evolutionary processes.

2.2.1 Groups

The notion of distributed evolution allows us, for the satisfaction of the observer, to define *groups*. If the distributed structure can be divided into two subsets such that the retention and the revival mechanism operate only within the two sets but not between them, we say that the two subsets form two groups. Each group is a structure undergoing its own evolutionary process.

Groups can emerge in several ways. One possibility is the physical isolation of the two groups (allopatric speciation). Given such isolation, there can be no interaction between the groups. Another possibility is that two groups become incompatible with each other (sympatric speciation). When this is the case, a group is not capable of interacting with other groups except indirectly, by being part of the other groups' environment.

We can also note that the definition of groups depends on the time frame of the observer. If an evolutionary process is observed during a sufficiently short period, no interactions will occur between the individuals, and hence, every individual constitutes one group. Thus, when one wants to distinguish groups in an evolutionary process, one must consider the time scale environment that the process inhabits.

2.2.2 Generations

Generations are traditionally included in evolutionary models, but they are optional and not required in every evolutionary process. The contribution of generations is only to offer a way of allowing modification of the preservative structure. The occurrence of generations means that individuals in the distributed structure are *replaced* by new individuals. The modification of the

preservative component, however, may be introduced by other means. The existing individuals can be *modified* instead of being replaced by new ones.

Associated with generations is the genetic *code*. In computer simulations of 'evolutionary processes,' some kind of genetic string is often used. The purpose of this string is to provide a way of implementing a variation mechanism. This mechanism, however, does not have to operate directly on the preservative structure (as the genetic code) but can act indirectly through the active structure component, as pointed out above. Genetic codes are thus not required in all evolution.

2.2.3 Fitness

Fitness is a measurement of how well a structure in an evolutionary process is adapted in relation to the environment. The fitness is implicitly given through the mechanism of selection. Structures that are selected are said to have a higher fitness than structures that vanish. The fitness is thus a value assigned by an observer *a posteriori* to selection, and is not a part of the evolutionary process itself.

In many computer simulations, however, the fitness is explicitly assigned to each structure and the selection mechanism is then based upon this measurement. Such a procedure assumes a meta-level, or an external observer able to evaluate the entities, and is thus not an example of evolution. At least, not in the sense that evolution is construed in this paper.

3 ONTOGENETIC ADAPTATION

In the presentation of the general framework above, we have already touched upon the issue of ontogenetic adaptation. We will now go into more detail concerning how this framework can be applied to the adaptation of the individual.

3.1 THE BRAIN AS A UNIVERSE

The first thing that we need to recognize is that the evolutionary process of ontogenetic adaptation takes place in the abstract space – the universe – that the brain forms, i.e., the relationships between the possible neural activity patterns that the synaptic configuration defines. These relationships form a structure, and it is this structure that is adapting.

As evolution takes place in the brain of an individual, the process is by necessity limited to the life time of the host. This has several implications. First, the process must run much faster than the evolution of the species. Otherwise, there would not be enough time for any adaptations to occur. Second, all the accumulated adaptation – the knowledge – of the process will be lost when the host organism dies. This is not true,




Figure 2: Distributed evolution. s_i , a_i and p_i represents individuals, while A and P represents solitary structure components. a) The structure is not divided into components. b) The active and the preservative components are distributed together. c) The active and the preservative components are individually distributed. d) The active component is distributed and the preservative component is solitary. e) The preservative component is distributed and the active component is solitary.

however, if we take communication between individuals and the evolution of culture into account, but I will not go into that in this paper.

3.2 COMPONENTS AND MECHANISMS

The relationship between possible neural activity patterns is, as mentioned, the adapting structure. The active component of this structure is the current neural activity pattern. The preservative component is the current synaptic configuration. Note that the structure components are abstract constructs and do not correspond to any physiological structures in the brain. Nevertheless, these components are distributed over physiological entities. The active component is distributed over the individual activity state of the neurons, and the preservative component is distributed over the state of each synapse (connection) between two neurons. Thus, the structure components of ontogenetic adaptation are distributed separately.

To see how the mechanism operates on these structure components, we start by noticing that whenever a neuron is active, synaptic modifications will follow. Hence, the activity pattern of the neurons will affect the synaptic configuration, or, in other words, the active structure component affects the preservative. This relation we recognize as the retention mechanism.

Next, the synaptic configuration will affect the way in which neural activity patterns may form. One may say that the current activity pattern creates the next such pattern by propagating activation through the connections. The updating of activity patterns will thus be influenced by the synaptic configuration. In

other words, the preservative structure component will affect the active. This relation we recognize as the revival mechanism.

So far, nothing has been accomplished but new names to well-known properties of neural networks. The revival mechanism is just the way activities are propagated, and the retention mechanism will be recognized by many readers as a Hebbian-like learning rule.³ The introduction of the remaining two mechanisms, selection and variation, requires a more novel view of the operation of neural networks.

The best way to introduce these mechanisms is perhaps by the use of an example. Let us consider the case of a motion detection network. This network consists of a lot of neurons, where each must emit an activation when a stimulus is moving in a certain direction at a certain position on the retinal image. Thus when an image is moving across the retina, we will be able to see the same image moving across the matrix of the neurons that responds to that specific direction. The parts of the retinal image that are not moving, or moving in another direction, will, however, not appear in that matrix.

Let us now see how we can implement this behaviour. To do this, I will ask you to think of yourself as taking the place of a motion detecting neuron

³Some readers may object that a Hebbian-like *learning* rule does not fit into an evolutionary process. Learning does, after all, impose a learner. This objection is, however, not justified. The Hebbian-like learning is taking place at a lower level of complexity than the structure that is evolving. That is, the Hebbian-like learning is learning about the relation between two neurons. It does not know anything beyond that. Hebbian-like learning is *not* learning at the global level where the evolving structure is situated.

with the mission of detecting leftward motion whenever it occurs in the small visual field that you are allowed to observe. What would you do? Well, my first guess is that you would cheat! Just think of the temptation. You are surrounded by neurons with a similar mission. If a pattern moves toward you in the specific direction that you are obliged to detect, a lot of neighbouring neurons will signal this motion before it reaches you. All that you have to do is to watch them carefully. When they are signalling motion, then you simply await for the pattern to appear in your visual field. Since you already know that it is moving in the proper direction you can, with confidence, signal an occurrence of motion and hope that no one will discover your fake.

A probable objection to this is that your trick will work fine as long as it is only one neuron that is cheating and *one* neuron will not have any considerable impact on the network's computation. Anyhow, I will, a bit contra-intuitively, state that we can construct a motion detecting network where *all* the neurons are cheaters. After all, the arrangement will work *as long as an initial detection is put into the system*. A pattern that is once detected, will be propagated by the cheating neurons until it disappears or halts. If the patterns halts, the neurons still get the information from their neighbours, but the pattern will never appear in their visual field, and, consequently, they will not signal.

The question that remains to be solved is how the initial detection is to be done. This is where the variation mechanism comes in. If we at every step randomly select a small percentage of the neurons to signal irrespective of any motion, they will falsely inform their neighbours that a moving pattern is approaching. What will happen is that *if* a pattern really is approaching, the cheaters will propagate the information and a detection of the moving pattern will take shape in the network. If the falsely induced signal is *not* in correspondence with any moving pattern, no pattern will appear in the cheaters' visual fields and hence, the 'detection' will not be propagated any further and disappear. Clearly, the detection that emerges in the network is under the influence of the retinal input (the environment). This influence we recognize as the selection mechanism.

The example above has been successfully implemented and tested using video recordings as stimuli (Pallbo, 1993; 1994a,b). Compared to other implementations of motion detection, this model is unusually insensitive to noise in the input. In addition, the architecture of the network is very simple. What makes this possible is the usage of spontaneous activity that builds up the perception. Such spontaneous activity is also found in most biological neural tissues (Evarts, 1964; Burns et al., 1976; see also Freeman and Skarda, 1990), but ironically, many neuroscientists view it as contributing nothing to the process (Adey, 1970).

It should be clear that the methodology is not limited to the domain of motion detection. The stimuli (i.e., the selectional pressure) imposed on the network can come from any area of the brain, not only the retina or some other perceptor area, but also from an internal source. Regardless of what domain the input originates from, the spontaneous activity will cause a neural activity pattern to emerge that is 'in resonance' with the imposed input. This activity will, in turn, affect the synapses by mean of the retention mechanism, so that the same pattern, or parts of it, will be revived more easily in the successive process. Furthermore, structural relationships in the imposed input that are not earlier recognized may be found and retained in the preservative structure.

3.3 KNOWING WITHOUT A KNOWER

When the evolving structure is exposed to the environment for the first time, the creation of neural activity patterns does not profit from any retained knowledge. It is the selectional pressure from the environment alone, but with the possible exception of the addition of innate preferences, that sculpts the neural activity patterns. For this reason, only the simplest, or most obvious, external structures of the environment will be captured. The fundamental environmental structures that are retained will then serve as a platform for the recognition of increasingly more complex structures in the following process. Thus, during the evolution of the structure, it will reflect more and more advanced characteristics of the environment to which it is exposed.

In a similar manner, the phylogenetic adaptation reflects the environmental conditions in the phenotype. For instance, the body of the fish is conditioned by the environment of the fish – i.e., water. In the same way, aerodynamics are reflected in the phenotype of birds. The phenotype does not have to know this. It is in and of itself a constitution of the knowledge. The body of the fish does not have to know that it constitutes knowledge about water, it is knowing without a knower.

Turning back to ontogenetic adaptation, we can make a parallel argument. A neural activity pattern stands, by (our) definition, in a certain relationship to other possible patterns. That is, the neural activity pattern has a certain place in the structure. From this position the pattern obtains its semantics. There would be no gain in introducing some entity that observed this fact, because it is already factual. That is, the neural activity pattern does not have to know that it constitutes this knowledge. It is, again, knowing without a knower. It is a knowing mind, not a mind with knowledge attached.

4 DISCUSSION

What we have done in his paper is basically to assume the existence of some kind of structure that is subject to change. This structure may be many things. To name a few, it may be a biological organism, an atom, a culture, or, what we have been concerned with here, a mind. The event of changing the structure has been referred to as the variation mechanism. This general scenario is in line with some other ideas. Pirsig divides the world into static and dynamic quality in his *Metaphysics of Quality* (1991). The static quality corresponds to what we have called structure in this paper, and the dynamic quality is the energy that causes change. In addition, the old philosophy of Tao holds a similar concept of the world. In this paper, however, I have adopted the evolutionary terminology. Furthermore, I have taken the risk of educing a series of special cases from the general setup. These special cases were derived from the idea of dividing the structure into two components, as well as the idea of distributed evolution.

When we considered the special case of ontogenetic adaptation, we identified the various mechanisms that are in operation. We must, however, not fall prey to the temptation of believing that we can *describe* these mechanisms other than at a meta level. If we make a description of how the mechanisms should act upon a specific structure, we have restricted the process to evolve within the frames of this description. Besides, it would not be an evolutionary process any more, but instructive. An evolutionary process has no place for a someone or something that produce descriptions of how the evolution should progress. In the terminology of the metaphysics of quality, we can say that the mechanisms, when described, have become static and no dynamic quality is to be found. The philosophy of Tao expresses this as “The tao that can be told is not the eternal Tao.”⁴ We can make a meta level description, and we can make an *a posteriori* description of how the mechanisms have acted upon a structure, but we can never give a corresponding *a priori* description.

4.1 RELATED WORK

The view of ontogenetic adaptation presented in this paper is not the first one to employ evolutionary terminology. In this paper, however, the model of ontogenetic adaptation as evolution is perhaps more actual than metaphorical as I feel it to be in related work.

To suggest *selection* as part of the ontogenetic adaptation is not a new idea. In 1967, Jerne remarks that many attempts to explain biological phenomena with instructive models have later been replaced by selective models. He invites us to ask ourselves “if learn-

ing by the central nervous system might not also be a selective process; i.e., perhaps learning is not learning either.” (p. 204).

To use selection in the adaptation process of a neural network can be viewed in contrast to traditional connectionism (e.g., Rumelhart and McClelland, 1986). The models within this community are based on the assumption that the architecture of the connections in the network constitutes a parallel machinery capable of computing various tasks. Along with the network comes a *learning algorithm* that is used to change the strength of the connections in the network and thus affects the computations that this network makes. Thus “[c]onnectionism shares with AI the underlying assumption that adaptive behavior is the result of computational processes” (Reeke Jr. et al., 1989, p. 136). The learning algorithm takes a position of a *knower* in the system. As a consequence, the connectionist system cannot adapt to situations that were not taken into account during the construction of the system, since these situations are simply not detected.

4.1.1 Neural Darwinism

The Theory of Neuronal Group Selection, also known as Neural Darwinism (Edelman, 1987), is perhaps the most well-known model of ontogenetic adaptation that uses selection. The basic idea is that the ontogenetic development of the brain supplies it with more neurons (clustered in groups) and connections than are needed. During the encounter with the world, the ‘best’ groups of neurons and connections are selected while the other degenerate (Edelman and Reeke Jr., 1982; Edelman, 1989; Reeke Jr. and Sporns, 1990; Edelman, 1992; see also Crick, 1989).

Neural Darwinism is thus based on an assumption of a repertoire of neuronal groups with a preexisting diversity. The neuronal repertoire must be sufficiently large to allow a match of each stimulus with a neuronal group in the repertoire. This view has the strange consequence that if the repertoire grows larger than ‘sufficient,’ no extra gain is achieved (Edelman, 1979). Furthermore, the idea that the diversity is preexisting is in conflict with what many biologists mean by Darwinism; “The dance evolutionary biologists call the ‘Darwinian two-step,’ randomness-then-selection continuing back and forth for many rounds to increasingly shape up nonrandom-looking results, usually cannot be seen in Edelman’s examples of neural Darwinism” (Calvin, 1988, p. 1802).

The model presented in this paper does not assume a ‘sufficiently large population’ of neurons. The entire population of neurons is used from the very beginning, and the more neurons in the repertoire, the larger the structures that can be kept in the system. In addition, the model proposed in this paper uses a variation mechanism that operates continuously on the structures.

⁴These are the first words of the ancient book “*Tao Te Ching*” by Lao Tzu. Here in translation by Mitchell 1990.

4.1.2 Pre-representations

A second approach of ontogenetic adaptation that involves selection is the model put forward by Changeux and his group (Changeux and Danchin, 1976; Changeux, 1983; Changeux and Deheane, 1989; see also Piattelli-Palmarini, 1980). This model views neural representations in the same way as presented in this paper, i.e., as the activity pattern of the neurons. Furthermore, this model also employs spontaneous activity to allow a mechanism of diversity. The spontaneous activity is used to generate ‘pre-representations’ in the brain. If a pre-representation is in resonance with a percept, that pre-representation will be selected. Once selected, it will form a ‘stored representation.’ This contrasts with instructive models where the storage is explicit.

Even if Changeux’s approach is similar in several respects to the approach taken in this paper, there are some crucial differences. In the model proposed here, the idea of pre-representations is not used. Rather, the representations are continuously constructed ‘on the fly.’ The variation and selection mechanisms operate at the level of individual neurons rather than on the level of neuronal patterns which has consequences for how acquisition of knowledge is accomplished.

Another important difference between the pre-representation model and the model of this paper, is the way in which the components are interpreted. Changeux et al. speak of the brain as a *user* of these representations (1984, p. 118). The model presented here, however, strives to eliminate any kind of user of the representations. The representations constitute their own structure and need no interpreter or knower.

4.2 WALLACE’S PROBLEM

Since ontogenetic adaptation by tradition has been viewed upon as a *learning* process, much of the research has been aimed at finding the structures (in the brain) that constitute the learner. These structures have been basically searched for in the physiological structures of the brain. This is perhaps why connectionism has grown to such popularity. Others, however, suggest that “something like software” is doing the learning (e.g., Dennett, 1991, p. 190). What both of these approaches have in common, though, is the assumption that the evolution of the brain as a physiological entity has supplied it with the necessary structures underlying a mind. This concerned Wallace as he felt that the brain is too sophisticated for the needs of our ancestors (Milner, 1990, p. 457). Why should evolution construct anything that is unnecessarily complex?

With the evolutionary view of ontogenetic adaptation, however, comes the view that what should be searched for, in the physiological structures and processes of the brain, is not the entity (distributed or not) doing the learning, but how these structures implement

a universe and what characteristics this universe has.

Maybe it was in the brains of the first mammals that a small universe first appeared – a universe where the ontogenetic adaptation could emerge as an evolutionary process. To host such a process would probably be very beneficial for the animal and add up to its evolutionary fitness. This trait would thereby be able to evolve further. Furthermore, the phylogenetic adaptation is consequently provided with a much simpler task. Rather than having to evolve a complex structure underlying the necessary components of the mind, it only has to do the much simpler task of expanding the universe in which the mind is spontaneously self-created. Perhaps it is this expansion of the mind’s universe that underlies the almost explosive evolution of the mammals – a ‘Big Bang’ of ontogenetic adaptation.

ACKNOWLEDGEMENTS

I thank Matthew Hurst for inspiring discussions and his recommendations of relevant references.

This work has been supported by the Swedish Council for Research in the Humanities and Social Sciences.

REFERENSER

- Adey, W. R. (1970). Spontaneous electrical brain rhythms accompanying learned responses. In Schmitt, F. O., editor, *The Neurosciences: Second Study Program*, pages 224–243. Rockefeller University Press.
- Burns, B. D., F.R.S., & Webb, A. C. (1976). The spontaneous activity of neurones in the cat’s cerebral cortex. *Proceedings of the Royal Society London (Biology)*, 194:211–223.
- Calvin, W. H. (1988). A global brain theory. *Science*, 240(June 24):1802–1803.
- Changeux, J.-P. (1983). *Neuronal Man*. Oxford University Press. Translation of *L’homme neuronal* (1983).
- Changeux, J.-P. & Danchin, A. (1976). Selective stabilization of developing synapses as a mechanism for the specifications of neural networks. *Nature*, 264:705–712.
- Changeux, J.-P. & Deheane, S. (1989). Neuronal models of cognitive functions. *Cognition*, 33:63–109.
- Changeux, J.-P., Heidmann, T., & Patte, P. (1984). Learning by selection. In Marler, P. & Terrace, H. S., editors, *The Biology of Learning*, pages 115–133. Springer-Verlag.
- Crick, F. (1989). Neural edelmanism. *Trends in Neuroscience*, 12(7):240–248.
- Dagani, R. (1992). Synthetic self-replicating molecules show more signs of life. *Science*, (February 24):21–23.

- Dennett, D. C. (1991). *Consciousness explained*. The Penguin Group.
- Edelman, G. M. (1979). Group selection and phasic reentrant signaling: A theory of higher brain function. In Schmitt, F. O. & Worden, F. G., editors, *The Neurosciences: Fourth Study Program*, pages 1115–1139. Rockefeller University Press.
- Edelman, G. M. (1987). *Neural Darwinism: The theory of Neuronal Group Selection*. Basic Books.
- Edelman, G. M. (1989). *The Remembered Present: A Biological Theory of Consciousness*. Basic Books.
- Edelman, G. M. (1992). *Bright air, brilliant fire: on the matter of the mind*. Basic Books.
- Edelman, G. M. & Reeke Jr., G. N. (1982). Selective networks capable of representative transformations, limited generalizations, and associative memory. *Proceedings of the National Academy of Sciences*, 79:2091–2095.
- Evarts, E. V. (1964). Temporal patterns of discharge of pyramidal tract neurons during sleep and waking in the monkey. *Journal of Neurophysiology*, 27:152–171.
- Freeman, W. J. & Skarda, C. A. (1990). Representations: Who needs them. In McGaugh, J., Weinberger, N. M., & Lynch, G., editors, *Brain Organization and Memory: Cells, Systems and Circuits*, pages 375–380. Oxford University Press.
- Havel, I. M. (1993). Artificial thought and emergent mind. In Bajcsy, R., editor, *IJCAI-93. Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 758–766. Morgan Kaufmann Publisher.
- Jerne, N. K. (1967). Antibodies and learning: Selection versus instruction. In Quarton, G. C., Melnechuk, T., & Schmitt, F. O., editors, *The Neurosciences: A Study Program*, pages 200–205. Rockefeller University Press.
- Menger, F. M., Eliseev, A. V., & Khanjin, N. A. (1994). “A self-replicating system”: New experimental data and a new mechanistic interpretation. *Journal of American Chemical Society*, 116(8):3613–3614.
- Milner, R., editor (1990). *The Encyclopedia of Evolution*. Facts on File, New York.
- Mitchell, S. (1990). *Lao-tzu: Tao Te Ching*. Kyle Cathie, London.
- Pallbo, R. (1993). Visual motion detection based on a cooperative neural network architecture. In Sandewall, E. & Jansson, C. G., editors, *Scandinavian conference on artificial intelligence – '93*, pages 193–201. IOS Press.
- Pallbo, R. (1994a). Motion detection – a neural model and its implementation. *LUCS Minor*, 1. ISSN 1104-1609.
- Pallbo, R. (1994b). Motion detection by cooperation – a case study of the time dimension in neural networks. Technical Report LU-CS-TR:94-125, LUNFD6/(NFCS-3086)/1–71/(1994), Lund University, Dept. of Computer Science.
- Piattelli-Palmarini, M., editor (1980). *Language and Learning: The debate between Jean Piaget and Noam Chomsky*, chapter 8, pages 184–202. Harvard University Press, Cambridge, Massachusetts.
- Pirsig, R. M. (1991). *Lila: An inquiry into morals*. Bantam Press.
- Plotkin, H. C. & Odling-Smee, F. J. (1981). A multiple-level model of evolution and its implications for sociobiology. *Behavioral and Brain Sciences*, (4):225–268.
- Reeke Jr., G. N. & Sporns, O. (1990). Selectionist models of perceptual and motor systems and implications for functionalist theories of brain function. *Physica D*, 42:347–364.
- Reeke Jr., G. N., Sporns, O., & Edelman, G. M. (1989). Synthetic neural modelling: Comparisons of population and connectionist approaches. *Connectionism in Perspective*, pages 113–139.
- Rumelhart, D. & McClelland, J. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1: Foundations*. MIT Press.
- Sober, E. (1993). *Philosophy of Biology*. Dimension of Philosophy Series. Oxford University Press.
- Tjivikua, T., Ballester, P., & Rebek, Jr., J. (1990). A self-replicating system. *Journal of American Chemical Society*, 112:1249–1250.
- Watts, A. (1975). *Tao: The Watercourse Way*. Pantheon Books, New York.