General Introduction

This is a journey into noise. Even though noise is frequently considered as a nuisance – a disturbance of the actual signal – we will in this text see how noise can be not only useful, but also *essential*, to cognitive processes. The models presented in this work depend to such an extent on the disturbance of the "signal" that they would not function without it. The disturbance that noise brings into the system allows new patterns of thought to emerge (cf. Bateson, 1967). In fact, we will see how such patterns may grow out of noise. Therefore, instead of a nuisance, noise becomes the primary motor of every cognitive process.

Such a noise oriented approach to cognition will lead us away from the computational view of the mind and the idea of cognition as the processing of information. Instead, cognition becomes a process of growth. As such, we will liken cognition to an evolutionary process. Evolution offers in fact a model of creative processes of growth in which noise is an essential ingredient. In contrast to biological evolution, however, cognition does not span over many generations of a species but is limited to the individual brain and the life-time of the individual. Due to this difference, the concept of biological evolution is not directly applicable to the activities that go on inside the brain. It is, however, possible to extract the crucial features of evolution in order to create a more generalized theory. Within the frames of such a generalization, it is possible to depict cognition as an evolutionary process. The first paper in this collection, "Mind as Evolution and Evolution as Such," outlines such a general theory and shows how it may be applied to the processes of the mind.

Once the mind has been brought into an evolutionary framework, it is possible to discuss the mind's ability to reveal structures hidden in the environment, i.e., its learning capabilities. The nature of this approach, however, puts certain demands on the concept of knowledge. The view taken here is that knowledge, engined by spontaneous bursts of neural activity, grows by itself. The classical view of knowledge as true or false beliefs is alien to this approach. Instead, Popper and Campbell's evolutionary epistemology is used as a basis for the knowledge concept (see e.g., Radnitzky and Bartley, 1987). The second paper, "An Inquiry into Meno's Dilemma," expands this issue. By a series of arguments it is demonstrated that evolution can be seen as a knowledge process and therefore, that the views of the mind as an evolutionary process and as a knowledge process do not exclude each other, but on the contrary, are mutually embracing. Furthermore, this paper puts great emphasis on the mechanism whereby knowledge grows, i.e., on the demands of the mechanism of variation in evolution.

To bring these theoretical constructions into more concrete terms, a model of visual motion detection, based on the theory, has been developed. Like the theoretical work, it addresses the ability of the mind to absorb structures that reside in the environment into the cognitive process. In this case, the structures in question are movements in the visual scene.

Most models of visual motion detection are founded on the paradigm of information processing. That is, visual motion detection is performed by a process that takes two consecutive images as input. The process then computes and correlates the differences of the images and finally produces the estimated motion as output. Such an approach tries to bring order to a chaotic visual scene by means of rigid algorithms . The approach taken here, on the other hand, based on the evolutionary framework, turns this completely around. Here, motion in the visual field is viewed as a *stable* phenomenon used to stabilize an otherwise chaotic brain. Only in the presence of movement do the "detectors" of motion in the brain display an ordered behaviour. In the absence of persistent movements, there is only a noisy background activity. Thus, instead of assuming the visual scene to be chaotic and the cognitive process to be ordered, *cognition* is assumed to be chaotic and the visual scene offers the stability. The last three papers of this thesis present a model that allows the detection to emerge, or grow, from the background noise. These papers also present the results of a computer simulation based on the model using video sequences of a real world scene as stimulus.

The elimination of a homunculus

What then are the benefits of viewing the mind as an evolutionary process? What does this approach offer that cannot be accounted for by more traditional models like computationalism or connectionism? Just like evolution eliminated the need of a creating agent, a god, from the explanation of the origin of the species, evolution can eliminate the need of a homunculus from the explanation of cognitive process.

In most models of cognition the operative energy is assigned to a homunculus. That is, it is a homunculus that makes things happen in the process – it is the source of initiative. We will speak of a homunculus whenever the operative energy is controlled by a knowledge process, that is, whenever the source of the operative energy contains knowledge about the domain in which it operates. The designation "homunculus" is justified because the process tries to "know" what it is doing. The other option in focus here is to use a *blind* mechanism. That is, a mechanism without knowledge of the domain in which it operates. Thus, the process cannot *perceive* events in that domain, nor interpret its own actions.

In evolution the operative energy is realized by the means of a variation mechanism. This is a blind mechanism that basically introduces a random disturbance to the evolving structure. The random disturbance can be realized in many ways. Usually, it is realized through a genetic mechanism, but in the evolutionary framework of the mind the process has no genes at its disposal. Therefore, the random disturbance in the mind is realized through the noisy background activity of the neurons, that is, by the activities of the neurons that take place spontaneously and sporadically without a "proper" stimulus. As neural activity is directly related to the evolving structure, this approach is very advantageous compared to proposals that make use of genetic mechanisms, i.e, to models that involve some "genetic" string that is unfolded into a phenotype. In contrast to genetical evolution, the evolving structure does not need to be expressed as a string. In this way, all the problems concerning coding and interpretation of strings are avoided. This is otherwise, in my opinion, a serious problem in simulated evolution. It happens very easily that a knowledge process slips unnoticed into an interpretation process, i.e., that the interpretation process contains knowledge about the code. Whenever that happens, a homunculus has been reintroduced in the process and the benefits of the evolutionary approach is lost. In simulated evolution this will happen if we let the genetical string evolve, but not the process that unfolds this string into a phenotypic structure (i.e., the ontogenetical process).

With some care, though, a homunculus is avoided as a precondition in evolution. Quite ironically, however, it is a homunculus that we wish to obtain, because that is how we experience ourselves. The second paper argues quite extensively that evolution may be viewed as a knowledge process. This view is necessary in order to explain how an evolutionary process can account for the mind.

If it sounds paradoxical that we at all costs want to avoid a homunculus, at the same time as we wish to obtain one, this is only because the jump in the level of organization is not being recognized. It is correct that we want to obtain a knowledge process *as a global result*, but we do not want it at the expense of introducing a homunculus as a component. If a component homunculus is controlling some part of the process, the resulting knowledge process can never grow beyond the knowledge contained in that component. Furthermore, we would also have had to explain the source of that knowledge – a far from trivial task that is avoided when no homuncular component is assumed.

When all the components of the resulting homunculus, i.e., the mind, are free from any knowledge of the domain in which the mind operates, it is also free from inert restrictions that dictate where and how to look for new knowledge. Instead of having the growth of the system determined by its components, it is determined by the environmental structures to which it is exposed. Thus, the environment has taken a crucial and active role in the cognitive process. This is very beneficial because the knowledge that we wish the mind to absorb already resides in the environment. By utilizing the environment in the cognitive process in this way there is no need for additional knowledge processes in the components.

Top-down, bottom-up, and hither and thither

It may be illustrative to ask whether the approach to the mind as an evolutionary process postulate a bottom-up or a top-down process.

A bottom-up process is a process that starts from the most elementary structures in the input. In a visual system these could correspond to the intensity at each point. From these elementary structures the process will educe more and more complex and compound structures, such as the edges and solid patterns in a visual system. Clearly, this process starts from the "bottom" and works its way "upward."

A top-down process, in contrast, starts with an idea or an expectation of a relatively complex nature. This idea or expectation is then used as a guide in the processing of input stimulus. That is, it may suggest the presence of certain structures in the stimulus which may be searched for to verify the expectation. A classical example of a top-down process is a person, afraid of the dark, who walks through a dusky forest. This person may suddenly perceive a beast looking out from between the trees. At a second look, however, she realizes that what appeared to be a living being was nothing but a dead trunk and some branches. The hallucination could arise because the person *expected* (or perhaps feared) to see such things. This expectation was then used to guide the perceptual process in a top-down manner.

The top-down and bottom-up approaches suggest a *direction* of the process which proceeds either downward or upward (or both) in a rather orderly fashion. The mind viewed as an evolutionary process fits neither of these approaches, nor a mixture of top-down and bottom-up processes. The evolutionary process does not start or end at any point and it proceeds in all kinds of directions. All activity originates from the background noise and grows in all directions from the point where it first struck root. Furthermore, there may be several such points of origin simultaneously and at different levels of organization. Therefore, it makes more sense to talk about a *hither and thither* process rather than a top-down or bottom-up process.

Situated cognition

The idea of letting the environment take part in the cognitive process is not new. The approach is called "situated cognition" which indicates that the cognitive process is always situated somewhere. It does not, or cannot, take place in empty space.

An advantage of situated cognition is that the cognitive agent does not need to be fully informed about the features of the environment which it inhabits. It only needs to know the environmental properties that are essential to the kind of behaviour that it needs to perform.

Much of insect behaviour can be modelled along this line. A tick, for instance, may sit in a tree or on a blade of grass and wait for a mammal to pass by and become a host animal to it. The tick, however, does not need to know about the existence of mammals. It does not even have to know that it sits in a tree or in the grass. All it needs to know is that whenever it detects the presence of a certain chemical compound, it should loosen its grip. If the tick is lucky, it will drop on the bypassing mammal who is emitting this very chemical. If the tick is not lucky, then it should reposition it-

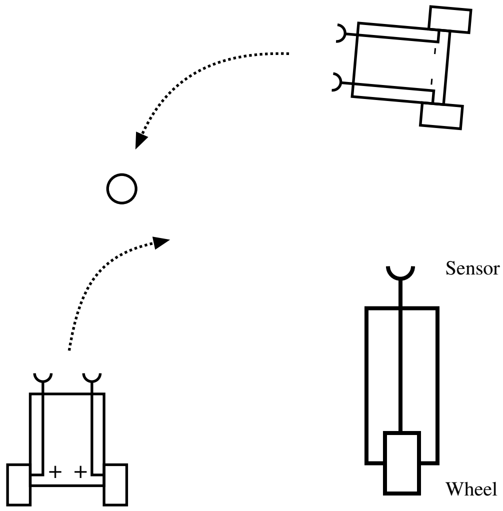

Figure 0.1: By connecting sensors to affectuators (or wheels) various behaviour can be accomplished. To the right we see the simplest possible vehicle with one single sensor connected to one wheel. The other two vehicles are composed of a pair of sensors and wheels which are connected so as to accomplish avoidance behaviour (bottom) and slow approaching behaviour (top) respectively.

self and give it another try.

This kind of behaviour may be synthesized. In a very delightful book, Braitenberg (1984) described a series of creatures, or "vehicles" as he prefers to call them, able to act in a simple environment. In the more simple exemplars of these vehicles, various behaviours arise by simply connecting sensor elements to affectuators (typically wheels which enables the vehicle to move around) in various ways (see Fig. 0.1). Even if the behaviours that these vehicles are able to perform are very simple, there is at least one thing in their favour – they work. Attempts using classical AI to construct robots, in contrast, are not renown for their success.

The reason for the success of these creatures lies mainly in their simplicity. Without complex processes of symbol manipulation, there is not much that can malfunction. This is not to say that only simple creatures can be constructed within the approach of situated cognition. They may be extended by more complex behaviours. Such an extension must be done wisely, however, or the organizational simplicity is lost. To enable wise extensions, Brooks developed something he called the *subsumption architecture* (1991a; 1991b).

In Brooks' architecture, simple, but fundamental, behaviours are first implemented and tested thoroughly. An example of such a simple behaviour implemented in a mobile robot may be to move away whenever an object is too close to it and to remain still and inactive otherwise. To this basic repertoire, gradually more complex behaviour may be subsumed. An explorative behaviour may, for instance, be added to the object avoidance behaviour mentioned above. That is, instead of remaining still, the robot can randomly move around in its surroundings. Whenever an obstacle comes too close, however, the basic avoidance behaviour is triggered and takes control. Once the obstacle has been avoided, the explorative behaviour may be continued. On top of this explorative behaviour, yet another level may be implemented, for instance a more goal-directed movement where the robot tries to move toward distant places.

Brooks has actually implemented this very set of behaviours, used in the example above, in a robot called "Allen" (1990). He has also constructed a set of other robots along the principles of the subsumption architecture. E.g., "Tom and Jerry," two reconstructed toy cars who exhibit a following behaviour when they move about in the same environment, or "Herbert" who picks up soda cans from office rooms. Brooks and his team at the AI lab at MIT have been encouraged to such an extent by these and other robots that they have initiated a program to build a robot in humanoid form (Brooks and Stein, 1993).

To let the environment take part in the cognitive process is indeed a powerful approach. In the approach taken in this work, where the mind is considered as an evolutionary process, the environment is employed not only in the cognitive process but also in the learning process. Throughout this work, this learning process is viewed as an evolutionary process in the individual mind. It might, however, be enriching to add yet another perspective to this process, namely the theory of autopoiesis.

Autopoiesis and structural coupling

The theory of autopoiesis was founded by Maturana and Varela as an attempt to define the characteristics of life. Instead of claiming that life is characterized by reproduction, breathing, or any other such distinct feature, they claimed that the essence of life is to be found in the way organisms are organized. A living organism, they say, in contrast to a dead one, is an organization that continuously produces and maintains its very own organization. It is a homeostatic process where the process itself is the variable that it tries to keep constant (Maturana and Varela, 1980). Such organizations they call autopoietic which literally means to create oneself.

To enable this self-creation, the properties of the medium in which the organization emerges becomes crucial (Fig. 0.2). It is crucial for the very same reason as why the realization of the variation mechanism in an evolutionary process is important. The medium cannot be allowed to be passive or no organization would take shape. The properties that allows the autopoietic organization to emerge must be properties already present in the components of the autopoietic entity. The properties cannot

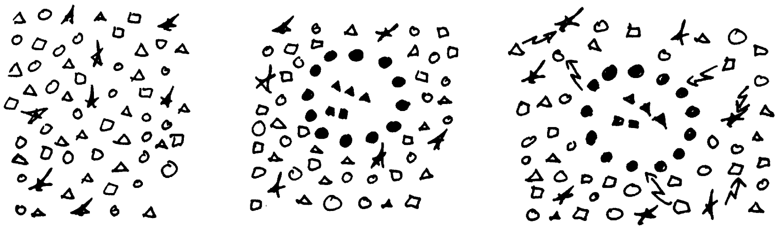

Figure 0.2: In autopoiesis one starts off from a medium (left) in which autopoietic organizations emerge (middle) by means of the relations between the substances in the medium (right).

be imposed from external sources because there are no external sources in autopoiesis.

This lack of external sources lead us to a central issue in autopoiesis, namely that of unity. According to Maturana and Varela's definitions, a unity is an entity that is separated from the background by an operation of distinction. In general, this operation of distinction is performed by an observer that is external to the system. In autopoiesis, however, this operation is performed by the autopoietic system itself.

This means that when life is treated as an autopoietic unity, the organism's existence is independent of any external observer recognizing it. This way, the autonomy of life can be accounted for. In contrast, a non-autopoietic unity such as a car, exists only as a unity when there are humans as external observers to recognize it as such. The car does not define its own existence. I believe that this difference between unities that do and do not define their own existence is crucial also in the study of conscious experience. To obtain a first person experience, one cannot rely on external observers. Therefore, it is necessary that a conscious unity defines its own existence. Autopoiesis provides the tools to describe such systems.

What constitutes the autopoietic unity is the organization of the components. That is, the relations between the components define the system as a unity. Maturana and Varela are careful to distinguish between this *organization* of an autopoietic system and its *structure*. While the organization refers to the instrumental participation of the components in realizing a unity, the structure refers to the actual nature of the components and their actual relations, which may include other than those necessary to realize autopoiesis. If we go back to our car used in the example above, the organization of the car consists of the relations between the components that allows the car to function as a car, e.g., the relations between engine and the wheels or between the steering wheel and the wheels. The structure of the car, on the other hand, refers to the physical realization of these components, e.g., to the *rubber* wheels, to the properties of the *steel* engine, etc.

With these definitions and distinctions at hand, we are ready to confront the phenomena in autopoiesis that are of most interest to the issue of this book, namely perturbations and structural coupling. As mentioned above, the medium – or the domain in which the structure exists – must be active. This activity allows the organization of an autopoietic system to emerge, but it also causes the system to be exposed to perturbations due to this very same activity. This is also the reason for why an autopoietic system not only must create itself but also maintain itself. When the organization is disturbed, the perturbation must trigger a reaction in the autopoietic system. If the system fails to compensate for a perturbation, the autopoietic organization may disintegrate and the unity is lost, or, in the case of a living organism, it dies.

An autopoietic system that successfully manages to maintain an invariant organization under the structural changes, that are generated by the active medium, has established a set of specific reactions that are triggered by specific perturbations. These reactions and perturbations result in a structural correspondence, or a *structural coupling*, between the changes of states in the medium and the changes of states in the autopoietic system that follows. From an external point of view – to an observer – this coupling will appear to be a *semantic* coupling (Maturana, 1975). But this semantic coupling is completely located in the eye of the observer. From the point of view of the system itself, the coupling is nothing but the result of the components' interactions with the medium. The observer, however, is able to *interpret* what it observes and may easily mistake the additional information to reside in the system. In a similar way, people are very inclined to attribute human properties to their pet animals or even to their cars.

The distinction between structural and semantic coupling is crucial. While structural coupling refers to the *structure* of the system, semantic coupling refers to the *organization*. Since the organization is a property of a specific unity, it requires an interpretational effort to recognize it. Therefore, a structural coupling is much easier to establish than a semantical coupling.

These two different viewpoints are of interest to the discussion of the concept of knowledge in the second paper, "An Inquiry into Meno's Dilemma." There it is claimed that evolution can be viewed as a knowledge process. Not all philosophers are comfortable with such a view, however. The reason for this discomfort is probably due to the fact that it corresponds quite well to the external view of an autopoietic system. That is, the knowledge process can be said to correspond to the apparent semantic coupling in autopoiesis. In the end, however, the structural and the apparent semantic couplings refer both to the very same process. It is just two views of the same thing. Likewise, the evolution and the knowledge process refer to the same process but from two different viewpoints. In working with these concepts it is vitally important to recognize this fact – not least when the issue of conscious experience is addressed.

Final remarks

In concluding this introduction, I would like to add that there is one aspect of cognition that is not covered by the models in this present work. It is the question of evaluation.

Suppose that an organism is engaged in some kind of activity. This activity will lead to transformation of the cognitive organization to allow the organism to remember and repeat the same behaviour again on a later occasion. This is also how the model is presumed to operate in its current state. If a certain pattern, or sequence, of neural activity has been present in the cognitive system, it will change the system in such a way that the same pattern, or sequence, will be more likely to re-occur. After many repetitions, the pattern or sequence will be even more likely to reappear.

But sometimes, it may be better to *avoid* a certain behaviour if it has proved to be harmful to the organism. If a behaviour in the past has lead to injury it could be a good strategy to avoid that behaviour in the future.

Evaluation is a very delicate issue that has to be handled with great care. We cannot introduce a process that directly evaluates a behaviour if we want to refrain from reintroducing a homunculus. But what we do have at our disposal is the emotional reactions of the organism. These reactions are of course also an evaluation, but it is not an evaluation of the specific behaviours but of the general *state* of the organism. Such an evaluation does not have to be dynamic in the sense that it needs to follow the development of increasingly more complex behaviours. Since it is only the state of the organism that is evaluated, this evaluation is independent of the complexity of the behaviour that generated that state. It *is* a knowledge process, but it is knowledge about a domain that does not change within the lifetime of the individual. Therefore, biological evolution can account for this process.

The emotional evaluation process can be used to treat good and bad behaviours differently. Since it is the synaptic transformations that occur after neural activity that modify the successive behaviour of the system, it is at this level the difference must manifest itself.

A possible way to accomplish the effect is to let various emotional states produce different kinds of chemical agents that are released in the brain to modulate the way in which synapses transform. To date, however, the model has not been brought to this level of analysis. Therefore, this adjustment is still at the level of speculation on behalf of the author and needs to be brought into accord with known neurophysiological data. A logical continuation of the project described in this work is therefore to investigate this low level of operation and modulation of synaptic transformation.

The lower levels of cognition cannot be studied in isolation, however. All levels of cognition interact in an exciting way (cf. Havel, 1995). It has been stated in this introduction that the cognitive process has no beginning or end. This applies also to the various levels involved. Cognition does not start or end at any specific level, but all levels are involved continuously. Even though a cognitive process may originate in noise at the neural level, the way in which the process grows forth from this noise involves many more levels. The relatively high level of neural activity, for instance, depends on the lower level of synaptic modifications, while the relatively low level of synaptic modifications, in turn, depends on the higher level of neural activity. I suspect that we can find this circular dependency between more or less all levels of cognition, and, therefore, the effects of noise at the neural level quickly manifests itself at many levels simultaneously.