

Inside Embodiment — What means Embodiment to Radical Constructivists?

B. Porr and F. Wörgötter

Bernd.Porr and worgott@cn.stir.ac.uk

+44 (0)1786-46-6378/6369

Computational Neuroscience

Department of Psychology

University of Stirling

FK9 4LA, Stirling

Scotland, UK

Abstract

The nervous system is operationally closed. It operates only in contact to itself. This astonishing claim has been made by Heinz von Foerster, one of the founders of radical constructivism. This work explores the consequences of his claim in the context of linear signal theory, embodiment and the creation of artifacts. In linear signal theory all transfer functions can be directly associated with the neural activity where also the environment is described by neural activity. This means that the environment is no objective entity but is described in terms of internal activity. We construct the worlds in our heads. The phenomenon of embodiment is interpreted here from the perspective of the nervous system, thus from the inner perspective. To identify inside and outside an organism must learn to identify the disturbances which are only in the environment. This can be done by anticipatory learning. Thus, embodiment is a process which emerges from the anticipation of disturbances. If one wants to design artifacts in the context of constructivism one has to obey that only quantity (activity) plays a role but no interpretation of activity. In addition one has to design an agent which performs input control and not output control.

Keywords: Embodiment, Cybernetics, Closed Loop, Self referentiality, Forward Model.

1 Introduction

One of the fathers of radical constructivism is Heinz von Foerster (Von Foerster, 1960). He proposed the second order cybernetics which emerged from classical cybernetics introduced by Wiener (1961) and his colleagues. In classical

cybernetics closed loop control systems are observed from the outside by an objective observer. All variables are observable and the system can be described in an objective way (Kalman, 1963). Von Foerster argued that objective observation is not possible since also the external observer operates as a closed loop system. Observation means simply that the observer integrates the observed system into his or her own loop.

In this paper we will also deal with an observer problem: is it possible for the organism itself to observe its own boundaries? Thus, we radically employ the perspective of the organism. More precisely we will radically employ the perspective of the nervous system. Therefore we will ask if the nervous system is able to distinguish between inside and outside¹. In more modern terms we ask if the nervous system is able to decide if it is *embodied* (Brooks, 1989; Pfeifer and Scheier, 1999) or not.

Mathematically we will stay in the field of linear² control theory and we will try to give this mathematical formalism a new meaning in the light of radical constructivism (Von Glasersfeld, 1984; Schmidt, 1987).

Having introduced the mathematical formalism makes it easy to switch over to artificial agents (“animats”) which can be described by this formalism. Consequently, we will finally give some guidelines how to apply the highly theoretical claims to more practical situations.

2 Reactive systems

In this section we will elaborate on reactive systems (Phillips, 2000). We define a reactive system as a system which only acts after a sensor event has occurred. In the context of control theory a reactive system is known as a feedback system.

We will, however, show that the interpretation of the mathematical formalism is different if we radically employ the perspective of the agent. This section will basically put von Foerster’s ideas (Von Foerster, 1985) into the context of linear control theory (D’Azzo, 1988). We will do this in two steps: we will first introduce reactive autonomous systems. Secondly, we will discuss the issue of embodiment

¹In this work we assume that the nervous system is already there and that it is operational in the sense that it generates stable reactions of the agent. Such stable reactions could have been established by (artificial or natural) selection (Bulinaria, 2003)

²We don’t make the claim, however, that everything could be modelled in a linear fashion in real world applications. Linear control theory only serves as an example. The claims below can be easily extended to non-linear theories if one obeys causality.

in the light of feedback control. It will be shown that from the perspective of the agent it will be difficult to distinguish between environment and agent.

2.1 Transfer functions and signals

Linear control systems are described by transfer functions and signals (D’Azzo, 1988). A transfer function transforms an input signal into an output signal. If we switch from the temporal domain to the Laplace space (Stewart, 1960) even this distinction is no longer needed.

In this section we will show how the transfer functions and signals can be interpreted from the organism’s point of view.




Figure 1: A simple self referential system. The transfer function H_0 transforms sensor signals X_0 into motor signals V . The transfer function P_0 transforms the motor signals V back into sensor signals X_0 .

The control diagram in Fig. 1 represents a minimal system with feedback. The symbol H_0 represents the transfer function which defines the agent. The transfer function P_0 is the transfer function of the environment. The role of the transfer function H_0 is to transform input signals into output signals. Here it may help to think of a sensor-motor transform ($V(s) = H_0(s)X_0(s)$). P_0 on the other hand defines the transfer function of the environment which performs the opposite ($X_0(s) = P_0(s)V(s)$).

Both P_0 and H_0 form a closed loop system with its own dynamics. Forming a closed loop is an important aspect or even the starting point for researchers whose work is related to radical constructivism (von Uexküll, 1926; Maturana and Varela, 1980; Ziemke, 2001). Such a closed loop system can be stable or unstable. Von Foerster (1985, p.79) like other cyberneticians demands that such a system must be stable. Stability in our simple linear system means that there is a desired state of the system and this state will be reached after a finite time. The actual desired state is defined by both transfer functions P_0 and H_0 . To achieve

stability usually negative feedback is introduced which is a counterforce which compensates the disturbance³.

The two transfer functions can be combined to one transfer function $G(s)$ by, for example, dividing both transfer functions by the environmental transfer function $P(s)$. Thus, the transfer function $G(s) = H_0(s)/P_0(s)$ represents both the environment and the organism. This seems to be weird as this makes organism and environment indistinguishable.

To make sense of $G(s)$ we have to interpret $H(s)$ and $P(s)$ in a different way. Identifying $H(s)$ and $P(s)$ with the agent and the environment respectively doesn't work any longer. It makes more sense, however, to *identify both H and P with the agent*. Given this interpretation $G(s)$ would also simply represent the organism.

Identifying both functions $P_0(s)$ and $H_0(s)$ with the agent is a reasonable explanation if we are radically employing the perspective of the agent. In particular if we are here employing the perspective of the nervous system. One of von Fo-




Figure 2: Self referentiality: Nerve cells (N) are interconnected by synapses (syn) and by the sensor- (S) motor- (M) loop. They also influence themselves with the help of hormones (NH). Reprinted with kind permission of Tom von Foerster and Rodney Clough who did the original artwork.

³Negative feedback shall not be interpreted as a form of evaluation like a punishment. It simply means that any deviation receives a counterforce which restores the desired state. If a room is too hot the central heating will go off. If the room is too cold the central heating will go on. The negative feedback involves only a sign inversion but no evaluation.

erster's famous statements is that the nervous system *only* transforms neuronal signals into neuronal signals (Von Foerster, 1985, p.48). Signals transmit only quantities but no quality. A light sensor transmits only intensity but not that it is light (von Uexküll, 1926). From the perspective of the nervous system also the environment only transforms neuronal signals into other neuronal signals: The *activity* of motor neurons only causes changes of the *activity* of sensor neurons (see Fig. 2, taken from Von Foerster (1985)). Thus, the environment is only a special synaptic "gap" from the perspective of the nervous system. Knowing that the nervous system only operates with neuronal signals we see that the transfer functions above represent simply neuronal activity and their transformations. This is trivially the case for the agent's transfer function H_0 but it is also the case for the environmental function P_0 .

Only an observer can distinguish between H_0 and P_0 . The observer, however, has difficulties to identify the right transfer functions. Therefore an observer might argue that such transfer functions are also not known to the agent. This, however, is only a problem of the observer. For the agent the transfer functions simply result from the neuronal activity which is transformed into itself – no matter where this actually takes place. For the agent it is simply his/her own personal neuronal activity. The world is in our head.

In summary the transfer functions $H(s)$ and $P(s)$ simply represent transformations of neuronal signals. From the perspective of the nervous system there is no distinction between inside and outside. This is one of the main arguments of radical constructivism.

Thus, embodiment in the naive sense of having a physical body with physical boundaries does not make sense from the perspective of the organism. In the context of radical constructivism physical embodiment does not exist⁴. We will, however, try to rescue the concept of embodiment by introducing disturbances from the environment.

⁴However, we won't argue against any physical instantiation of the nervous system which is interacting with its environment. Therefore a form of "classical embodiment" is involved. One might be tempted to compare our definition of embodiment with the standard definitions of embodiment (Chrisley and Ziemke, 2002). These standard definitions of embodiment, however, are difficult to compare to our definition: for example, the nervous system might not even perceive itself as "physically embodied" because its feedback loops include the environment. On the other hand we need some aspects of "organismoid embodiment" or even "organismal embodiment" to establish working feedback loops. It seems so that our definition is orthogonal to the classical definitions which arises from the different perspectives: external perspective vs internal perspective.

2.2 Disturbances

At this point we refer to von Foerster's colleague Ashby (1956) who introduced disturbances⁵ (denoted by D here). A disturbance is a signal which causes a deviation from the desired state in the feedback system. Every feedback system experiences disturbances. We get, for example, hungry which can be interpreted as a deviation from the desired state. To eliminate this undesired state we usually eat to restore the desired state. An optimally designed feedback system is able to cope with any disturbance. In Ashby's terms: The system has enough requisite variety to cope with any disturbance. In biology, however, this might not be necessary: McFarland (1989) argues that for an organism it is sufficient to maintain a "weak" homeostasis which keeps the organism away from lethal boundaries but allows in general deviations from the desired state.

What makes a disturbance different from an ordinary (neuronal) signal? A disturbance can also be formulated by a transfer function $D(s)$. From a pure mathematical point of view D is not special. The disturbance, however, adds the *unpredictable* component to the control system. It can not be absorbed by the control loop.

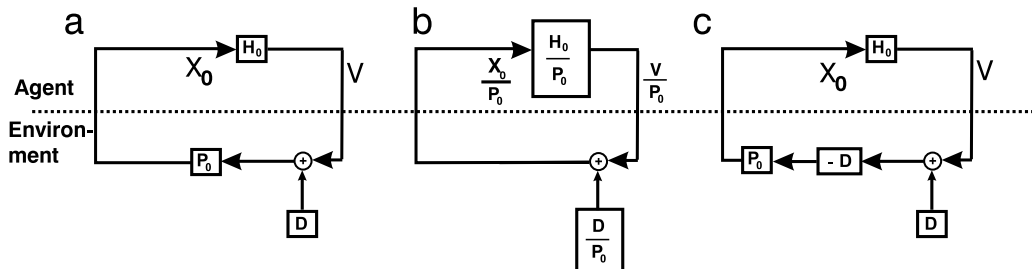


Figure 3: Transformation of the standard feedback-loop (a) into a unity-gain feedback (b). The transfer-function of the environment P_0 can be integrated in the transfer-function of the organism. However, any unpredictable disturbance can not be eliminated. c) A predictable disturbance, on the other hand, can be subtracted.

Now we have to define more precisely what actually a disturbance is. We have to add the disturbance to Fig. 1. This results to Fig. 3. Anything which is *predictable* cannot be a disturbance. Or, technically speaking, if we know or if we

⁵Maturana called the disturbances "perturbations" to make clear that he stresses the inner perspective of the organism. We will stay with the word disturbance.

can predict the properties of D , we can subtract it by including the negative transfer function $-D$ into the loop (see Fig. 3C). For example a constant (unchanging) input can be fully compensated for in this way and this holds for any other fully predictable signal. Indeed also this happens in many situations in physiology. For example, if a person permanently loses one side of his/her vestibular sensor input, he/she will feel this as a strong disturbance only for a while. After several days the system adapts to this permanent change of its inputs and the person fully recovers. Also periodic inputs are no disturbances. They also can be integrated in the loop.

Thus, it is the aspect of *contingency* which constitutes a disturbance and the existence of such disturbances is the constituting necessity for the existence of negative feedback loops. Without disturbances such loops would be unnecessary.

This makes it also clear that disturbances only make sense *for the agent* but not for the environment. Disturbances disturb the homeostasis of the agent. The environment, however, does not maintain homeostasis. Therefore, only the agent knows what is a disturbance and what is not. The environment can not be disturbed. The environment is not a closed loop system itself. Again, it must be stressed that this is the perspective from the agent. In the environment, however, there might be other closed loop systems which maintain homeostasis. From the perspective of the agent, however, these systems (organisms, ecosystems, . . .) are still interpreted as “the environment”:

One must distinguish “the environment” from systems within the environment. The environment contains many more or less complex systems, which have contact with the system for which they are environment because it is part of *their* environments and therefore an object of possible operations (Luhmann, 1995, p.182).

In general, it cannot be assumed that the environment is disturbed by the agent. This would demand an objective external observer which is not given in second order cybernetics. Consequently a definition of embodiment, as suggested by Quick and Dautenhahn (1999), which assumes that the environment is perturbed cannot be used in the context of second order cybernetics.

Now we have to ask the question: has anything changed? Is the agent able to distinguish between inside and outside if we make the disturbance explicit? The answer is yes and no. Yes, because the disturbance can be assumed of being outside. No, because the organism is still not able to *identify* the disturbance. The disturbance could have entered the loop at any place.

Consequently the next question which has to be addressed is: How can an agent identify the disturbance in the outside world? This leads to the general question: what is the agent able to observe? Thus, we have to postpone the question if the agent is able to distinguish between inside and outside. We first have to give an answer to the question what observation means from the perspective of an agent.

3 Input Control

This section will show that organisms can only observe their inputs but never their own outputs. This claim will be derived from the more general concept of second order cybernetics which argues that every observer is part of a closed sensor-motor loop.




Figure 4: The organism as an observer: *H* transfers a sensor-signal to a motor-reaction. *P* is the property of the environment and transfers a motor-reaction into a sensor-stimulus. The organism as an observer is only interested in those aspects of its own behaviour which feed back to its sensor-inputs (a). Any behaviour which never feeds back can not be of any interest (b).

Let us interpret the input control in the light of the findings of the last sections. We have learned that the nervous system only operates in self contact. This holds also true for the contacts with the environment. Now we can think how (motor) actions are actually observed by the organism. From the organism's point of view only actions which *feed back* to the organism's *sensors* can be observed (see Fig. 4). Any other action which simply disappears in the environment cannot be observed by the organism. Thus, there is no other chance for the organism as to analyse its inputs as this is the only aspect that the organism is able to observe. Even its own actions are only observable through its inputs.

What does this mean for embodiment? We have seen that the boundaries from the perspective of the organism no longer exist. The concept of input control, however, brings us back on track towards a more abstract form of embodiment. Input control is strongly related to the disturbances we've introduced in the last section. Disturbances appear at the sensor inputs of the organism and trigger a compensation reaction which in turn is again evaluated at the sensor inputs of the organism. We have argued, however, that this disturbance can enter the loop at any stretch of the loop. Thus, no distinction between organism and environment is possible so far. A solution of the problem arises if we take another sensor input which gives us an explicit information about the disturbance itself. If such an input exists the disturbance as such is identified. If we assume that the organism is not generating disturbances by itself it can conclude that the disturbance has originated in its environment. Thus, if we are able to gain information about disturbances we gain information about the environment.

One might argue that even the simple feedback system is able to distinguish between inside and outside. This however, is not the case. The simple feedback system only has to cope with the disturbance but it needs no knowledge where the disturbance enters the loop. This is the power of any feedback system. It needs least knowledge about its environment. One could say: it needs no explicit knowledge about its environment.

The main advantage of the feedback system is, however, also its curse: It can only react after a disturbance has happened. In von Foerster's words: this is the blind spot of the feedback system (Von Foerster, 1985, p.26). To make this clear we have to recall that the input X_0 measures *deviations* from the desired state. It measures, however, not the disturbance itself. This is due to the fact that the disturbance triggers a, often complex, reaction of the feedback system. Thus, the feedback system H_0, P_0 can not distinguish between cause and effect. Again, only the external observer sees that there has been a disturbance which has entered the feedback loop. The organism itself is not able to determine if the *deviation* from the desired state was due to the disturbance itself or due to its own reactions caused by the disturbance. In more mathematical terms: The input X_0 receives a *sum* of the disturbance and the output of the organism. Thus, the organism cannot distinguish between the disturbance and its own reactions against the disturbance.

Despite of these problems there is a way for the organism to detect the disturbance in its environment. The organism has to find a way to separate the disturbance from its own reactions against it. Thus, the organism needs another input which is not influenced by its own reactions and which carries the disturbance as a signal.

4 Anticipatory systems

In this section we will show that the agents are finally able to distinguish between inside and outside. This also opens the field to semantics because now “objects” in the outside world can be labelled. Let us emphasise again that we will focus on the inner perspective of the agent, not considering any external observation (and evaluation) of behaviour.

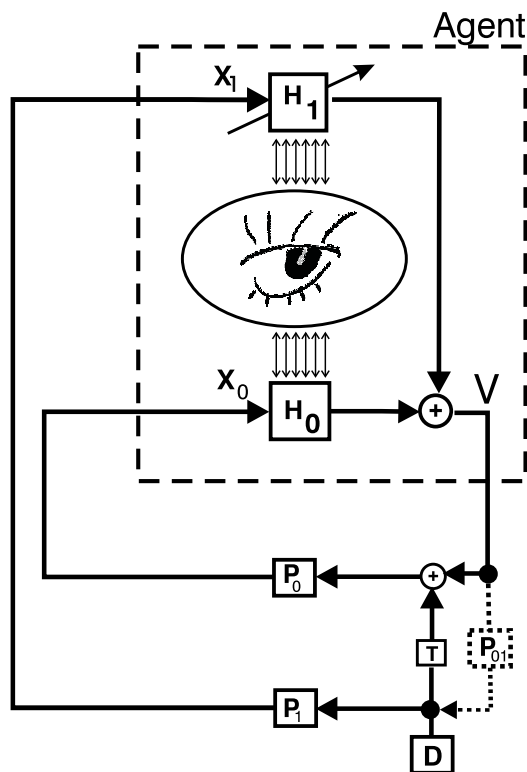


Figure 5: The inner feedback loop is observed by the outer feedback-loop. The inner feedback loop is established by the transfer functions H_0 and P_0 . The outer feedback loop is established by the transfer functions H_1 , P_{01} , P_1 . D is the disturbance and T delays the disturbance. The eye stands for the observation process: the outer feedback loop observes the inner feedback loop and adjusts H_1 so that the inner reflex loop is no longer needed.

Now we extend Fig. 3a by adding another input to the organism (see Fig. 5). When the agent has come to life only the innermost loop H_0 , P_0 is operational.

This loop has a desired state, which, however, can not be maintained all the time as disturbances (D) arrive at the loop occasionally. These disturbances D enter the inner loop delayed by time T . The *undelayed* disturbances enter the organism via the sensor input X_1 . In other words: The organism receives a predictive signal at the input X_1 in relation to the input X_0 . The signal at X_1 can now be used to observe the primary feedback loop (H_0, P_0) and determine what is the effect of D on the primary feedback loop. This observation can be used to adjust H_1 in a way that the inner feedback loop does not feel the disturbance X_0 any more.

How this works in practise can be seen in the following example (Porr and Wörgötter, 2003b): Imagine an agent is equipped with touch- and vision-sensors. It is moving around in a world which consists of rigid obstacles. The initial wiring of the agent is as follows: The primary feedback loop is established by its touch sensors which are connected to the motor outputs. The moment the agent receives a touch input (X_0) the agent moves backwards and turns. Thus, the desired state of the agent is to keep the touch sensor silent. This however, is not achievable as collisions will sometimes occur (D). There is, however, a way to avoid the trigger of the collision sensor. The signal from the vision sensor *predicts* the trigger of the touch sensors. The temporal correlation between vision- and touch-signals can be learned by means of temporal sequence learning (a more detailed description can be found in Porr and Wörgötter (2003a) and Porr, von Ferber and Wörgötter (2003)). After learning the agent produces an active avoidance movement before touching an obstacle. The same idea can also be applied to attraction cases.

What is learned through such a process? To understand this we realise that the reflex loop is for the agent an objective frame of reference and by this a *temporal* symmetry breaking point is defined: Measured against the *moment* the reflex is triggered, there are signals and actions which are “earlier” and other which are “later”. Thus, the reflex defines what is future and what is past. The desired state of the agent is to keep its touch sensor signals zero by avoiding the touch reflex. Thus, signals, such as the vision signals used here, which arrive *before* the reflex-eliciting touch-signal can potentially *improve* maintenance of the desired state. Improvement (or deterioration) can be *objectively* measured by the agent against the moment of triggering the reflex-loop. The outer loop implicitly attributes therefore a meaning to a reaction (e.g. “earlier” or “later”) with respect to the temporal frame of reference given by the inner (reflex) loop.

In the last section we demanded that we need another signal to identify the disturbance in the outside world. We have to recall that it was not possible for the feedback loop itself to decide where the disturbance has entered the loop. Having the additional sensor input X_1 the system can identify the disturbance

itself through the pathway $D \rightarrow P_1$. In Spencer Brown's words the system has introduced another distinction in the form of the input X_1 (Spencer Brown, 1969).

The problem however is still not completely solved if the transfer function P_{01} is not zero. This transfer function creates a new feedback loop which itself incorporates the problems of the inner loop. With $P_{01} \neq 0$ also the outer loop cannot identify the disturbance. Consequently a third and finally an N th sensor input is needed which is then finally not part of a loop. Therefore this is the chance to identify the disturbance and therefore the outside. Thus, having a nested feedback loop or a subsumption architecture makes the organism step by step "aware" of its outside. The more feedback loops arise the more the environment becomes distinguishable.

How many nested feedback loops can emerge? Theoretically as many as there are sensor inputs which predict each other. In practice, the predictions of the predictions will become more and more unreliable. Therefore learning will fail if the input signals become too uncorrelated to each other. Another aspect is the form of preprocessing. Predictions of predictions can be learned best if a simple sequence of events can be generated. Thus, if a complex sensor input can be converted to a simple sequence of signals predictive learning is simple to achieve. A typical example from biology are place cells which can be found in rats (O'Keefe, 1976; Blum and Abbott, 1996).

Let us leave this academic question and ask what the agent has learned in terms of signal and control theory. To clarify what the agent has learned about the environment we have to recall that the basis for learning was the fixed reflex reaction. In a living organism this fixed connection can be identified by a preprogrammed genetic disposition. The learning goal is to avoid this genetically defined reflex. This means that after learning the reflex-loop is no longer triggered. Thus, also the input X_0 is no longer needed. The organism detects a disturbance at X_1 and then issues a reaction which eliminates the disturbance before it can reach X_0 . In engineering terms this is called a forward model which has been learned by the organism. Also Rosen (1991) concluded that the organism calculates a forward model. Consequently the forward model which has been learned by the agent is the forward model of the primary reflex. In the case of nested feedback loops we get forward models of forward models. This reminds pretty much of Bateson's theory of learning (Bateson, 1979).

In the former section we have demanded that organisms perform input control. In this section we can extend this claim: Organisms perform *adaptive* anticipatory input control. Before learning the organism is able to cope with unpredictable situations because of its inherited requisite variety. After learning the organism

has learned how to use predictive sensor information to eliminate disturbances before they trigger its slow feedback reaction.

Having all this we can try two definitions of embodiment:

1. An organism is embodied if it has enough requisite variety to react against disturbances from the environment so that it stays away from its lethal boundaries.
2. An organism is able to distinguish between its inside and its outside if it is able to learn forward models of its own reflex-loops.

The first definition dates back to Ashby, has been modified by Mc Farland and has been adopted by Maturana. This definition means that if even the reactive system doesn't work the organism will disintegrate and die. The second definition dates partially back to Rosen who also demands that feedback systems have to develop forward models. The difference is, however, that Rosen observed the control system from outside as an external observer. The external perspective is questionable as it demands that the observer has a complete knowledge of the transfer functions. This however, is usually not the case (Dennett, 1984). Only the internal perspective eliminates this problem as pointed out above.

5 Practical consequences for the design of animats

In this section we are going to give some guidelines for constructivists who want to design animats (see also Ziemke, 2001).

5.1 The internal perspective

The internal perspective demands that we have to develop the animat from its own point of view (von Uexküll, 1926; Brooks, 1991). This seems to be trivial. This, however, is not the case. Let us recall the simple task of obstacle avoidance. We have to define obstacle avoidance from the perspective of the agent. The internal perspective demands that only quantities are correlated with quantities. This perspective forbids the integration of qualities into the design of the system. Consequently we have to establish rules which relate only signals with signals. In the case of obstacle avoidance this can be established by simple rules. Braitenberg (1984) has shown this very impressively by his thought experiments. The rules

only relate sensor signals to motor signals and the robot performs obstacle avoidance. They don't have to "know" that they perform obstacle avoidance. They only have to know how to transform signals. This is equivalent to an instrument-flight on a plane at night. The pilot reads instruments and operates handles. The pilot doesn't need to know that he or she is flying a plane (Maturana and Varela, 1980).

To make the design from "inside" clearer we shortly describe how a design from the "outside" might be developed. The external observer (for example an engineer) will probably start with the environment of the animat which surrounds the agent. Consequently he/she will start with objects in the environment. The objects form a certain pattern in the environment and therefore it seems to be straightforward to introduce a representation of the objects in the agent. Consequently an object recognition system must be built into the agent and the agent must develop a map of its environment. In fact many people draw this conclusion while watching Braitenberg vehicles and often attribute human features into the vehicles. Thus, the external perspective leads probably to completely different design goals. In particular the external perspective implies output control whereas the inner perspective implies input control. The consequences of input control will be discussed in the next section.

5.2 Input control

The inner perspective demands input control in contrast to output control (Von Glasersfeld, 1996). This means that the success or the failure of the agent's actions has to be measured at its inputs (sensors) and not at its outputs. If one wants to employ learning rules this has to be taken into account. Many neural network rules usually target a certain *output condition* and not a certain input condition like, for example, rules derived from the delta rule (Widrow and Hoff, 1960). A suitable learning rule is ISO-learning (Porr and Wörgötter, 2003a) which is based on differential Hebbian learning (Kosco, 1986). ISO learning stops if a certain input condition is met. Rules from the class of reinforcement learning (Sutton, 1988) might also be useful if they can be modified in a way that they target a certain input condition. In principle this seems to be possible⁶. To our knowledge, however, this has not been done so far.

At this point it must be noted that input control can lead to surprising results for the external observer: Since the agent controls its inputs the output can become quite unpredictable for an external observer. This, however, reflects only

⁶Jürgen Schmidhuber claims that this is possible. Personal communication.

the different intentions of an external observer and of an agent. Luhmann (1995) claims that this partial unpredictability of input controlled agents is the driving force for social systems where unpredictability is reciprocally generated between social agents (“double contingency problem”).

5.3 Anticipation

We have learned that anticipation is an important factor in autonomous learning (Rosen, 1991). The learning goal is defined simply by the “physics” of a reactive system, namely that it always reacts too late. Learning rules which tackle this problem stay on the level of signals (quantities).

Input control demands that an agent has to learn from its input statistics. This claim has been derived from the self-referentiality of the nervous system. Anticipatory learning in particular is based on the causality of events. The agent can not look into the future. It can only look into its past. Signal and control theory cannot look into the future. It can only look into its past. Consequently the favourite mathematical tool is signal- and control-theory in the Laplace space. Operations in the Laplace space automatically obey causality. They take into account the past but not the future.

Reinforcement learning also belongs to the class of anticipatory learning rules (Balkenius and Morén, 1998). As pointed out above, however, those rules must be adjusted in a way that they perform input control and not output control. Thus, in particular the reinforcement signal must be generated in a way that certain input conditions will be reached. This implies that the reinforcement signal must be generated by the nervous system itself so that it serves the internal goals of the nervous system and not the goals of the observer.

6 Conclusions

Von Foerster has given very important contributions to the theory of constructivism. They have been the basis of this article. In particular, his claim that the nervous system operates as a closed loop system has provided a definition of embodiment from the perspective of the organism. Embodiment from the perspective of the nervous system is not the starting point. It is rather the final result after learning. Before learning the closed loop structure of the nervous system does not allow a distinction between inside and outside. After learning the nervous system is able to detect disturbances in the environment which allows the nervous sys-

tem to identify inside and outside. The main result, however, is that the nervous system has learned to cope with disturbances from the environment. An observer might interpret this as gain of “competence” in the environment (Riegler, 2002) or as the “recognition” of objects in its world (Scheier and Lambrinos, 1996).

These more abstract results can also be transformed into practical guidelines for the development of artificial agents. If one wants to develop animats in the context of constructivism one has to start from the perspective of the animat. This means that only signals with other signals are correlated. This also means that the animat performs input control and not output control. Learning rules have to obey this. They have to change the internal structure according to the animats’ input statistics.

References

- Ashby, W. R. (1956), *An introduction to cybernetics*, Methnen+Co LTD, London.
- Balkenius, C. and Morén, J. (1998), Computational models of classical conditioning: A comparative study, Technical report, Lund University Cognitive Studies 62, Lund. ISSN 1101-8453.
- Bateson, G. (1979), *Mind and Nature — A Necessary Unity*, Wildwood House, Glasgow.
- Blum, K. I. and Abbott, L. (1996), “A model of spatial map formation in the hippocampus of the rat”, *Neural Comp.*, vol. 8, pp. 85–93.
- Braitenberg, V. (1984), *Vehicles: Experiments in Synthetic Psychology*, Bradford, Colorado.
- Brooks, R. A. (1989), “How to build complete creatures rather than isolated cognitive simulators” in VanLehn, K. (ed.), *Architectures for Intelligence*, Erlbaum, Hillsdale, NJ, pp. 225–239.
- Brooks, R. A. (1991), “Intelligence without representation”, *Artificial Intelligence*, vol. 47, pp. 139–159.
- Bulinaria, J. A. (2003), “From biological models to the evolution of robot control systems”, *Phil. Trans. R. Soc. Lond. A*, vol. 361, pp. 2079–2421.

- Chrisley, R. L. and Ziemke, T. (2002), *Encyclopedia of Cognitive Science*, Macmillan Publishers, chapter Embodiment, pp. 1102–1108.
- D’Azzo, J. J. (1988), *Linear Control System analysis and design*, Mc Graw, New York.
- Dennett, D. C. (1984), “Cognitive wheels: The frame problem of ai” in Hookway, C. (ed.), *Minds, machines, and evolution*, Cambridge University Press, Cambridge, pp. 129–151.
- Kalman, R. E. (1963), “Mathematical description of linear dynamical systems”, *J-SOC-IND-APPL-MATH-A-ON-CONTROL*, vol. 1, pp. 152–192.
- Kosco, B. (1986), Differential hebbian learning, in Denker, J. S. (ed.), *Neural Networks for computing: Snowbird*, Utah, Vol. 151 of *AIP conference proceedings*, American Institute of Physics, New York, pp. 277–282.
- Luhmann, N. (1995), *Social Systems*, Stanford University Press, Stanford, California.
- Maturana, H. and Varela, F. J. (1980), *Autopoiesis and cognition: the realization of the living*, Reidel, Dordrecht.
- McFarland, D. J. (1989), *Problems of Animal Behaviour*, Longman Scientific & Technical, Harlow.
- O’Keefe, J. (1976), “Place units in the hippocampus of the freely moving rat”, *Exp Neurol*, vol. 51, pp. 78–109.
- Pfeifer, R. and Scheier, C. (1999), *Understanding Intelligence*, MIT Press, Cambridge, MA.
- Phillips, C. L. (2000), *Feedback control systems*, Prentice-Hall International (UK), London.
- Porr, B., von Ferber, C. and Wörgötter, F. (2003), “Iso-learning approximates a solution to the inverse-controller problem in an unsupervised behavioural paradigm”, *Neural Computation*, vol. 15, pp. 865–884.
- Porr, B. and Wörgötter, F. (2003a), “Isotropic Sequence Order learning”, *Neural Computation*, vol. 15, pp. 831–864.

- Porr, B. and Wörgötter, F. (2003*b*), “Isotropic sequence order learning in a closed loop behavioural system”, *Roy. Soc. Phil. Trans. Mathematical, Physical & Engineering Sciences*, vol. 361 no. 1811, pp. 2225–2244.
- Quick, T. and Dautenhahn, K. (1999), Making embodiment measurable, in Proceedings of the 4th Fachtagung der Gesellschaft für Kognitionswissenschaft, Bielefeld, Germany.
- Riegler, A. (2002), “When is a cognitive system embodied?”, *Cogn. Syst. Res.*, vol. 3, pp. 339–348.
- Rosen, R. (1991), *Life itself: a comprehensive inquiry into the nature, origin, and fabrication of life*, Columbia University Press, New York, Oxford.
- Scheier, C. and Lambrinos, D. (1996), Categorization in a real-world agent using haptic exploration and active perception, in Maes, P., Mataric, M., Meyer, J.-A., Pollack, J. and Wilson, S. (eds.), SAB96 From Animals to Animats, Cape Cod, MA, September 9-13, MIT-press, Cambridge, MA, pp. 65–74.
- Schmidt, S. J. (ed.) (1987), *Der Diskurs des radikalen Konstruktivismus*, Suhrkamp, Frankfurt a.M.
- Spencer Brown, G. (1969), *Laws of Form*, George Allen und Unwin.
- Stewart, J. L. (1960), *Fundamentals of Signal Theory*, Mc Graw-Hill, New York.
- Sutton, R. (1988), “Learning to predict by method of temporal differences”, *Machine Learning*, vol. 3 no. 1, pp. 9–44.
- Von Foerster, H. (1960), “On self-organizing systems and their environments” in Yovits, M. and Cameron, S. (eds.), *Self-Organizing Systems*, Pergamon Press, London, pp. 31–50.
- Von Foerster, H. (1985), *Sicht und Einsicht: Versuche zu einer operativen Erkenntnistheorie*, Vieweg, Braunschweig.
- Von Glasersfeld, E. (1984), “An introduction to radical constructivism” in Watzlawick, P. (ed.), *The Invented Reality*, W. W. Norton, pp. 17–40.
- Von Glasersfeld, E. (1996), “Learning and adaptation in constructivism” in Smith, L. (ed.), *Critical Readings on Piaget*, Routledge, London and New York, pp. 22–27.

- von Uexküll, B. J. J. (1926), *Theoretical biology*, Kegan Paul, Trubner, London.
- Widrow, G. and Hoff, M. (1960), “Adaptive switching circuits”, *IRE WESCON Convention Record*, vol. 4, pp. 96–104.
- Wiener, N. (1961), *Cybernetics — or control and communication in the animal and the machine*, 2 edn, The M.I.T. Press, Cambridge, Massachusetts.
- Ziemke, T. (2001), “The construction of ‘reality’ in the robot”, *Foundations of Science*, vol. 6 no. 1, pp. 163–233.