“Inside” Embodiment – Control from the Organism’s Point of View

B. Porr and F. Wörgötter

bp1 and faw1@cn.stir.ac.uk
+44 (0)1786-46-6378/6369
Computational Neuroscience
Department of Psychology
University of Stirling
FK9 4LA, Stirling
Scotland, UK

Abstract

“Intelligent, flexible, and autonomous behaviour will only occur in embodied agents which are rooted in their environment.” This statement captures a currently widely accepted view about the requirements for designing advanced robots. Implicit to this statement is that un-embodied, pure algorithmic structures cannot become intelligent. In this article we will derive the minimal conditions for an embodied agent employing the view-point of systems theory. We will strictly employ the inner perspective of the agent, building our arguments only on its intrinsic signals (on which the agent has to exclusively rely) but not on its externally observable behaviour. We will show that negative feedback interactions of the agent with its environment are fundamental in establishing weak homeostasis by which the agent maintains its integrity. Furthermore we find that the border between agent and environment is solely defined through unpredictable contingent events called disturbances arriving from the environment at the agent. In control theory, events are in general described by transfer functions and it can be shown that it is possible to integrate every predictable transfer function *into* the agent. Thus, it is the environmental contingency which is the constituting element subdividing the agent’s *body* from its outside. Examples from animals and animats will be used to clarify the abstract concepts. As a result we arrive at the conclusion that computer programs can be embodied. Such agents are still pure syntactic, symbol-manipulation systems but we will show that they can develop very simple inner semantic evaluations by means of anticipatory learning which could be taken as a very first step towards intelligence.

Keywords: Embodiment, Intelligence, Observer Problems, Closed Loop, Self reference.

1 Introduction

Over the last years the discussion about what makes an agent intelligent has been augmented by introducing the so-called embodiment-principle [4, 17, 6, 7]. According to this, only embodied agents (artificial or natural) can perform intelligent interactions with their environment and during the last years the literature on this topic has become truly embodied [21, 27, 45, 8, 29]. Therefore the question arises what we have gained with the introduction of embodiment? Is it an equivalent to intelligence? Is it a necessary or a sufficient condition? And, can it be used to quantify or "calculate" any of these properties?

For a reader who is not familiar with the field of robotics or computer-science it may seem awkward that this community apparently needs to discuss the role of a body. It seems so obvious that without a body intelligence can not exist. Why does the field of robotics discuss the necessity of a physical body?

In order to clarify the motivation for the embodiment principle one needs to go back and look what had happened before embodiment was introduced. About 30 years ago the field of artificial intelligence experienced great fervour and hopes that artificial reasoning, expert systems and other approaches originating from this field could be a model for (human) intelligence. These hopes probably originated in the Turing-principle which essentially states that algorithmic computation is not linked to a specific substrate like, for example nerve cells or electronic circuits. Intelligence was taken to emerge from abstract and therefore disembodied computation.

The design of such classical AI-systems is based on the definition of symbols and operations which manipulate them. The symbols and the operations are defined by the designer of the AI-system. Typical AI-systems would, thus, try to implement a model of (some sub-aspects of) the world and perform symbolic operations on the modeled entities which it contained.

Alas, rather sooner it became clear that all these AI-systems were exceedingly limited in their performance: Typically these systems failed as soon as they were faced with real-world problems which would take them outside their sheltered artificial world-model. One of central problem which causes the failure of AI in the real world is the frame problem [10]: How much information does the algorithm need to cope with all eventualities from the environment? This is an unsolvable task for the (external) designer as he/she usually can not predict all future situations. The environment it too unpredictable to include all possible environmental influences into the agent by the external designer.

Another problem is the symbol grounding problem [11] which again relates to the external designer of the algorithm: Who gives the symbols meaning in the context of the algorithm? As a consequence of mainly these two problems, specifically in the field of robotics, researchers observed over and over again that a nice robot simulation would get into deep troubles as soon as it was implemented on a real robot.

At some point the following conclusion was drawn: Algorithms are *not* independent of their implementation. They are linked to a *body*. Thus, it was claimed that intelligence requires a (physical) body [4].

However, it is quite obvious that a body by itself is not enough for expressing intelligent behaviour. For example, a living and a dead fish have both a body and therefore appear both to be embodied. However, somehow there seems to be a striking difference between them. . . [8]. Therefore a body itself is not a sufficient condition for intelligent behaviour, however, it could be a necessary condition.

Central to this article is that we are going to explore embodiment from the *organism's point of view*. Like Robert Rosen [30] we will start with the organism as an entity which is protected by a boundary [35]. However, this boundary has to be permeable [3, 30] to a certain degree to allow for interaction with its environment. This permeability could be used for a first definition of embodiment like the one suggested by Quick & Dautenhahn [27] however we will suggest other definitions later on. A permeable boundary leads to the danger of disintegration. In order to prevent this, agents have to operate in a state of homeostasis ensured by a negative feedback loop through the environment. We will argue that, from the organism's view-point, feedback is the most fundamental form of interaction with its environment [38] creating a specific form of quantifiable situatedness [44, 2].

At the end of this article we will augment the aspect of feedback and we will argue, like McFarland [16, 17], that intelligence is linked to learning, especially learning of anticipations. Or more specifically: If an organism is able to learn to anticipate feedback reactions, then it has performed one of the first steps towards intelligent behaviour which conforms to [16].

In this article we will to use a systems-theoretical approach which has been developed in the field of control theory [36, 18, 9]. It offers the advantage that it is accessible to a rigorous mathematical treatment. This shall not be attempted here, but see [24]. It must be noted that control theory is not limited to active control. It also includes passive control. For example, a marble running into a pit can also be described by control theory. For a detailed description of the differences see [15].

2 A system's theoretical description of a body

In this section we want to develop embodiment from the perspective of the agent. We are going to develop four conditions which are necessary for embodiment.

2.1 Boundary

Defining a body seems to be easy by just introducing a boundary:

> A body is defined by a *boundary* which subdivides its inside from its outside[1].

This definition, however, gives rise to some naive assumptions, for example, in associating the boundary with a physical boundary (e.g. a membrane or skin). We would like to emphasise early on that this association is most of the time inappropriate, if not outrightly wrong. Boundaries must be seen in a much broader context which goes beyond a physical interpretation. For example, also a society can generate boundaries by forming groups (for example scientists which discuss embodiment) or themes [13]. In order to go beyond the naive interpretation of a boundary as a physical object one has to ask for the boundary's *function*.

2.2 Disturbance and Desired state

The function of the boundary is to protect the agent from *disturbances* coming from the environment [2]. This demands stability of the boundary over some certain time in the sense that it must maintain its function for the agent in order to keep it alive[2]. Note that the boundary is allowed to change in order to adjust to new/unknown disturbances from the environment[3]. However, the boundary must always be able to guarantee its function, namely maintaining stability.

Using the word "disturbance" implies that there is something which *can be* disturbed, namely a "desired state". We assume that it is the primary (implicit) goal of the agent to maintain its desired state in spite of disturbances. Therefore the function of boundaries can be expressed more explicitly:

[1]Note that the boundary is an *asymmetric* distinction defining a *limited* space inside and a possibly *infinite* space outside (see [35] for a set-theoretical treatment of such boundaries).

[2]In Maturana's words the *function* of the boundary is the *organisation* of a living being [14]. Organisation must not be misunderstood by Maturana's definition of structure. The organisation remains the *same* as long as the organism is alive.

[3]Maturana calls this "structural coupling" with the environment.

> The boundary has the function to maintain a *desired state* as good as possible by compensating against disturbances originating from the environment.

This means that disturbances *always* come from the environment. At some point later we will come back to the concept of disturbances which we believe is indeed the constituting element needed to define a body (see below).

However, it is not a realistic assumption that a boundary can maintain a desired state always and with infinite accuracy. Realistically every organism has tolerances for deviations from desired states [2]. Animals get attacked and injured, but recover later. One gets hungry and then eats thereby restoring the desired state of being sated. A deviation from the desired state is permissible as long as it is within certain (non-lethal) ranges. We will call this *weak homeostasis*.

Interestingly this concept goes beyond mere qualitative statements, because Ashby [2] has shown that quantification of the performance of a boundary can be directly obtained from the demand that a desired state has to be kept within a certain permitted range. Consequently such a quantification measures the deviation from the desired state given the strength of disturbances coming from the environment.

Also a clear definition of "death" can now be derived: If an agent (animal or robot) is no longer able to maintain weak homeostasis it is dead.

It is important to note, however, what is a disturbance and what is not can only be decided from *inside* the boundaries (body). An observer outside can never be sure if a stimulus is a disturbance, a desired input or if it even goes unnoticed. As a consequence, external observation leaves room for speculations (and misinterpretations) about the causality of events which drives an organism. Really only the organism "decides" what are disturbances which need to be eliminated in order to re-establish weak homeostasis and maintain integrity.

2.3 Permeability

Ultimate stability could be obtained by a totally impermeable boundary which would be inert against any disturbance from the outside. This, however, abolishes any kind of interaction of the agent with its environment. We could call it a *static stability*. This is, however, fundamentally incommensurable with intelligence, because intelligence is not a state, but a process and as such it can only manifest itself through some kind of interaction of the inside with the outside. As a consequence we have two assumptions which work to some degree against each other:

(2.2) Stability of the desired state and (2.3) exchange between the inside and the outside[4]. Summarising we can state that in general agents are open and closed at the same time. Open in the sense that they interact with the environment and closed in the sense that they need to employ weak homeostasis [3].

2.4 Feedback

This leads to the consequence that a process must exist which governs exchange and which at the same time guarantees stability. A multitude of possible processes exist (e.g., in engineered systems) which would assure this. Here, we argue that:

> Negative feedback is the process which most fundamentally assures stability against disturbances while it at the same time permits interactions of the agent with its environment.

Note, that the feedback principle is not an arbitrary assumption made by us. Instead, viewed from the perspective of the agent, is it truly fundamental, because from its own perspective *only those* actions that are feeding back to its sensors can be of any relevance to it. Actions that have an effect only on the rest of the world will not be perceived by the agent and are effectively non-existent for it.

Since from the agent's perspective every relevant motor action re-enters the agent through its sensors the entire feedback-loop can be described entirely by the agents internal signals (for example, neuronal signals as suggested by von Foerster [38]). Such a description can be interpreted as being sub-symbolic as assumed by Brooks [6] in his robot experiments. The central advantage of this view is that it allows using *control-theory* to describe the whole loop in a quantitative way.

We emphasise here on negative feedback. Negative feedback counteracts against disturbances and is therefore able to reestablish homeostasis. Homeostasis established by negative feedback is only "weak" homeostasis as negative feedback only reacts after disturbances have happened. Deviations from the desired states are permitted and even essential in order to guarantee interaction with the environment. Therefore there is no need for "perfect solidarity" between the organism and its environment [34] having a pure feedback control. During ontogenesis a system can achieve better (as judged by itself) maintenance of its desired state(s)

[4]Note that energy is required to maintain the boundary because the boundary is permeable. As a consequence there will be a (thermodynamic) tendency to smooth out any existing gradient. This can only be prevented by energy input into the system. The problem how to provide energy for the system has been extensively addressed by Prigogine [26] and Bertalanffy [3] and shall not be of further interest here.

by self-adaptation leading to improved homeostasis as we will show below with an example.

2.5 Summary

In this section we have defined an absolute minimally necessary structure (Fig. 1 A) that cannot be further simplified for an agent which could become intelligent: a boundary, disturbances, a desired state, permeability and feedback. These are the necessary conditions and we are going to describe the more abstract concepts by control theory [36, 18, 9, 19, 22, 20].

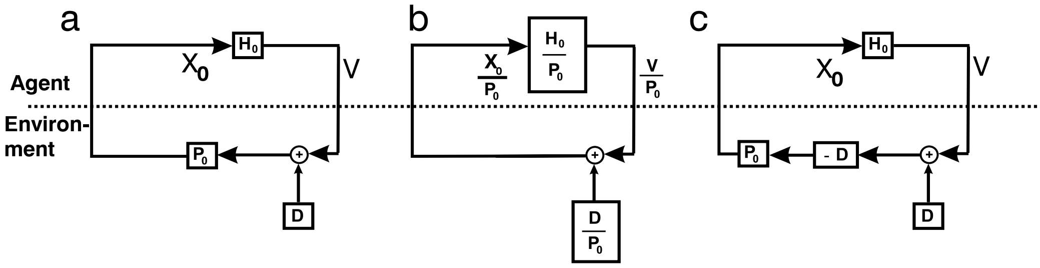

Figure 1: Transformation of the standard feedback-loop (a) into a unity-gain feedback (b). The transfer-function of the environment P_0 can be integrated in the transfer-function of the organism. However, any unpredictable disturbance can not be eliminated. c) A predictable disturbance, on the other hand, can be subtracted.

2.6 A minimal system with feedback

The control diagram Fig. 1A represents a minimal system with feedback. The symbol H_0 represents the transfer function which in the first instance defines the agent (or the body with its boundary, assumption 2.1), and by which it transforms its input(s) into some output(s). Here it may help to think of a sensor-motor transform. P_0 on the other hand defines the transfer function of the environment which (passively) performs the opposite.

For simplicity we assume that the desired state of the agent is given by $X_0 = 0$ such that its output should negatively compensate any disturbance D which occurs in the environment (negative feedback, assumption 2.4). If we want to have other desired states than zero one can always subtract the desired state from the reflex input x_0 so that the input of the transfer function H_0 becomes zero in the case of equilibrium.

The resulting loop represents the interaction of the agent with its environment (assumption 2.3). The fact that we have introduced negative feedback leads to the simplest form of weak homeostasis and, as a consequence, to stability (assumption 2.2). Note that there is no need to establish such feedback as active control in the form of electrically driven sensors and effectors. Passive, especially mechanical control can also be considered.

Such a system is strictly reactive, because it can only re-establish homeostasis *after* a disturbance has already taken place. Furthermore it operates purely in a syntactic way. This means that processing in such a system takes place by manipulation of signals and nothing but signals. Thus, we do not assume (or demand) that semantics ("meaning") and interpretations would exist in such a simple system. Really such a system is a simple feedback controller or a purely reflex driven animal. As a consequence such a system is not intelligent. How intelligence could arise, building on top of such a structure will be discussed later.

2.7 What does this mean "in everyday life"? - Exemplifying the transfer functions

In the following paragraphs we want to find examples what H_0, P_0 and D actually represent *for an agent*. To this end we will start with a technical observation stating that the diagram in Fig. 1A is problematic because one can integrate the transfer function of the environment into that of the agent (Fig. 1B). This is common practise in control engineering in order to simplify circuit diagrams. Now this seems like an awkward trick used to confuse the situation by creating a rather soft (changeable) boundary between a body and its outside. Suddenly we have created a "new agent" $\frac{H_0}{P_0}$ by combining H_0 and P_0 and this seems to be fairly counterintuitive because the distinction between environment and agent becomes very fuzzy this way. After all, we humans normally believe that our body ends at the skin beyond of which clearly there is the environment. This, however, is not necessarily always appropriate in practice. Let us introduce a living and a technological example thereby creating several apparently paradoxical "bodies" in oder to clarify this concept. Thereby we will at the end show that exactly this softing of the boundary takes place in the everyday life of humans (and robots).

First we consider the nervous system as such. Its properties define H_0^α and everything else is defined as the nervous system's outside P_0^α. This subsumes the sensors, muscles, bones, etc. of the physical body which surrounds the nervous system but also its outer surroundings given by the world. For example if this nervous system belongs to a human, we could consider the case that it sits in a car and drives on a road, which is clearly part of the outside, or is it not? The fact that the nervous system is comparatively stable over prolonged periods of time shows that it exists in some kind of (here unspecified) process that leads to homeostasis with its outside. Also there exists an interaction between the nervous system and the outside world and it will react to disturbances trying to re-establish its homeostasis. Thus, such a system obeys the structure of Fig. 1A. In this sense the nervous system would be an embodied agent. This looks queer. It seems so much more natural to assume that sensors, muscles, bones, etc. belong to this embodied agent and indeed within the framework introduced here, nothing prevents us from treating them as such. Thus, if we take some parts of the original P_0^α and combine it with H_0^α in the way suggested in Fig. 1B we arrive at a new set of transfer functions H_0^β and P_0^β. This way we have arrived at the (very abstract) description of the body of a human (or an animal). Interactions between P^β and H^β and the desire to maintain homeostasis also characterises this system.

Let us now come back to the example above and place this human into a car. Much in the same way as above we could now continue to extend the body of our agent towards the surface of the car by integrating the sensors, motors and the structural components ("bones") of the car into the body of the agent creating again a new set of transfer functions H^γ and P^γ. Indeed many drivers report that they regard the car as an extension of their body (this observation is even more pronounced in fighter plane pilots, when they report that they become "merged with the machine"). Also such a system interacts with its outside and tries to maintain homeostasis (e.g. in avoiding accidents). Common Sense, however, is very reluctant to accept such a seemingly arbitrary extension of a living body into a human/machine hybrid.

The same "Common Sense", however, does not seem to have any problem with this concept as soon as we transfer our example of a human/car "hybrid body" entirely into the non-living world. The nervous system could be associated with the control circuit of a modular robot. The human body-parts now turn into the metal, plastic and inanimate parts of the robot and there is no-one who would prevent us from designing a first-order robot (with some kind of physical body) which cannot drive and a second order robot (with an augmented physical body) which can drive. This way we have arrived at different embodied robots with different properties.

In the example of the robot we can even give up the integrity of the first order body when designing the second order body. Maybe this is the difference between animates and in-animates? Animates need to preserve their bodily integrity, whereas in-animates don't. A re-structuring of an agent will normally have to go through a state where homeostasis cannot be maintained (the robot is disassembled). At the moment this cannot be done with the physical body of a human. Disintegration and re-assembly still remains science fiction of the "beaming" process in Star-Trek.

However, even humans can be "re-assembled" to some degree (usually after accidents) and hybrids created between flesh and metal by integrating implants into the human. For example, many patients report that in the case of a limb-prosthesis this is soon felt to belong to their body. Furthermore, sensor implants, like an artificial cochlea or an artificial pacemaker for the heart are during normal operation never felt as being outside the patient's body because the patient does not have the means of self-observation as long as homeostasis is maintained. Only if the device fails, its property of being an integral body part fails as well and the device turns into a disturbance. The same, however, holds true for our original body parts as everyone who ever had tooth-ache can tell. Removal of such disturbances is highly desired.

While in the example of prosthesis, artificial body parts become integrated into the concept of the owner's body, the opposite is also observed. Cases are reported where a patient suffers from the loss of the mental image of some parts of his/her body ("negative ghost", [31]). The patient reports that the respective limb is not anymore felt as belonging to his/her body. The inverse situation is also observed; amputation patients often report that they still feel their missing real limb ("positive ghost").

All these examples show that from the organism's perspective boundaries can be drawn differently. Changing the boundaries means that the transfer-functions are changing like in the transition from Fig. 1A to B. In this case the entire environment is perceived of being part of the body. As mentioned above, this is common practise in control engineering in order to simplify circuit diagrams but it is also the central thesis of Constructivism how we construct our own world. The world is in our head – like radical constructivists, for example von Foerster [39], use to say. If we strictly adopt the perspective of the agent, who will have to operate only by its intrinsic signals, then there is no way for it to distinguish between different boundary-settings. The agent cannot find out where the influence of P_0 came from (the inside or the outside). Therefore, from the perspective of the agent P_0 could be an integral part of its body (its inside). This view is much more radical than, for example, the view by Weng [42] who demands that the organism must be aware of its own signals. We state that there is no other chance than of being self-aware of it own internal states as there are only internal states from the perspective of the organism. The organism can not distinguish between internal and external signals. Only an observer can perform this distinction. We will elaborate on observer problems later in this text.

Thus, we observe that the abstract concept of being allowed to integrate (some aspects of) the environmental transfer characteristic into "the body" (or doing the opposite) as depicted in the transition from Fig. 1A to 1B (or back) does indeed take place in the reality of human lives. It seems that even for a human the notion what constitutes his/her body is fuzzy under certain circumstances. Thus, the same should apply to robots.

As a consequence, our central conclusion is that everything which belongs to the feedback loop cannot be used to distinguish between the inside and the outside of a body. The only thing that objectively can normally never be integrated into the agent is the disturbance D. The disturbance (of the homeostasis) is always a property of "the outside". Ultimately only this can be used to subdivide inside (body) from outside (environment).

But what *is* a disturbance D? In the example of the aching tooth we have shown that an integral body part can also turn into a disturbance. So something which was definitely part of the loop (i.e.; part of H_0/P_0) now leaves the loop and becomes part of the outside. Thus, it seems that not even even disturbances are a unique distinguishing property in the temporal existence (lifetime) of an agent. Have we now reached a point where the embodiment concept has finally failed?

In answer to this question one has also to consider the temporal properties of such systems. Anything which is *predictable* cannot be a disturbance. Or, technically speaking, if we know or if we can predict the properties of D, we can subtract it by including the negative transfer function $-D$ into the loop (Fig. 1C). 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 looses 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. Pain adaptation is a similar example. Constant pain is in many cases to some degree ignored. This also happens with some forms of tooth-ache. This shows that human (or animal) bodies can indeed integrate predictable signals into their homeostatic loop(s). In the same way, unpredictable signals even when they arise from a formerly integral body part of the agent (like the starting tooth-ache) will be regarded as a disturbance and the agent will try to re-establish the desired state of the loop (e.g., “no pain”) by means of compensating for this new disturbance (drawing the tooth).

Thus, it is truly the aspect of *contingency* which constitutes a disturbance and the existence of such disturbances is the constituting necessity for the existence of negative feedback closed loops. Without disturbances such loops would be unnecessary. The existence of contingencies has been suspected to be the central underlying driving force of evolution as well as for the forming of social subgroups [13].

2.8 Can computer programs be embodied?

This extraordinarily broad concept of embodiment raises the question if it still has any explanatory power. Specifically one has to ask if the notion cited above that computer programs always fail to be embodied (and, thus can never be intelligent) would really hold. We believe [29] that computer programs can be embodied in our sense.

> The central constituting property of an embodied agent in the current framework is its homeostatic loop and its boundary with respect to contingent disturbances outside.

Computer programs can fulfil these two properties without any problem. The simulated agent shown in Fig. 2 is a practical implementation of Fig. 1a. It has a temporally stable boundary given by its numerically defined margins (assumption 2.2). It possesses a simple simulated reflex avoidance reaction in response to a simulated touch signal. Whenever it hits a simulated obstacle, it will perform a simulated retraction reflex through negative feedback (assumption 2.4). Thus, its desired state of homeostasis is to keep its touch-input signal equal to zero (see also assumption 2.2). This way it reacts to disturbances and thereby it does indeed interact with its simulated environment by being able to receive (touch-)inputs and perform simulated motor reactions (assumption 2.3). Thus, this agent is only a computer program and certainly it is not intelligent, but its visualisation in Fig. 2 is not necessary for its embodied (in our sense) existence inside the electronic circuits of its host computer.

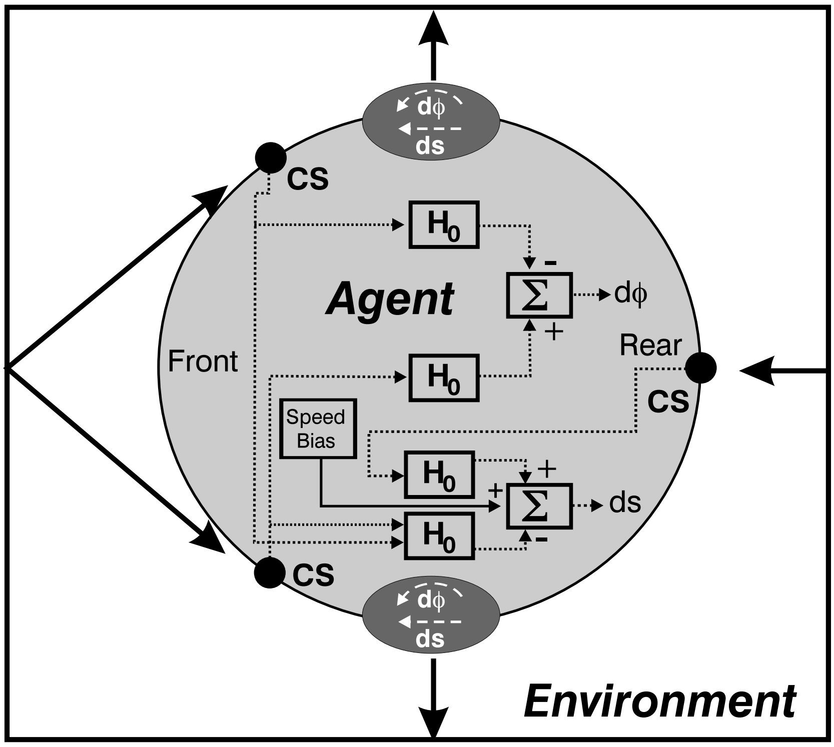

Figure 2: A reactive agent based on Fig. 1a. The agent is equipped with three collision sensors (CS) and two motors which are controlled by two neurons which determine the speed (ds) and the angle ($d\phi$). The transfer functions H_0 have the task to transform a sensor-input into a timed motor-reaction. The Speed Bias causes the robot to drive forward with constant speed as long as there are no bumps detected.

3 Towards intelligence

But is this very simple type of embodiment enough as a basis on which intelligent, flexible, meaningful (i.e. semantic) behaviour could be built, or did we just miss some important, fundamental aspects? There is no way to prove or disprove this rigorously given the fact that many viewpoints exist about what intelligence is. We can, however, show that such systems can be extended in a generic way and that by this extension a pure syntactic system turns into a system which can perform very simple (implicit) semantic evaluations. Let us emphasise again that we will focus on the inner perspective of the agent, not considering any external observation (and evaluation) of behaviour. A special section will be devoted to the problems of the "observer perspective" and that of "behavioural observations" later. Furthermore, one should keep in mind that the obtained augmented system is at the end still exceedingly simple such that we cannot claim to have achieved any kind of intelligence. The development of an inner semantic can only be regarded as a very first step towards this direction.

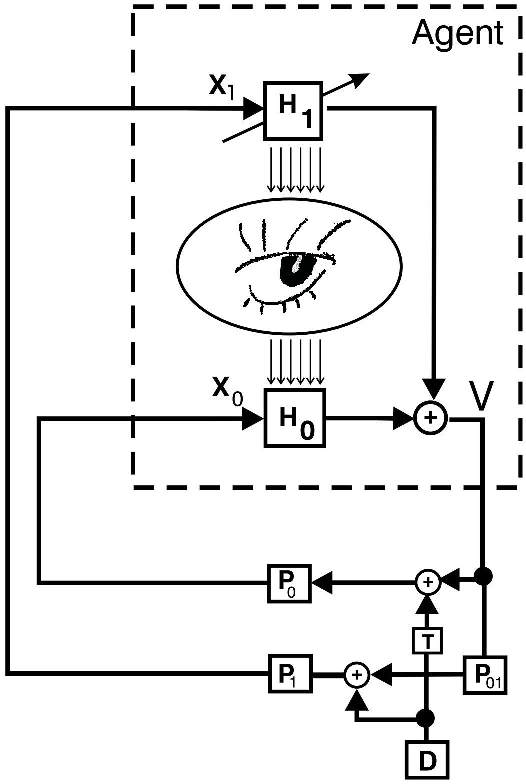

Figure 3: The inner feedback loop is observed by the outer feedback-loop. The inner feedback loop is established by the transfer functions H_0 and P_0. The outer feedback loop is established by the transfer functions H_1, P_{01}, P_1. D is the disturbance and T delays the disturbance. The eye stands 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.

In order to generate personal semantic evaluations we have to extend our very simple embodied agent from above (Fig. 2). This agent is reflex based (reactive) and its structure is that of a purely syntactic system. Now we design a so-called subsumption architecture of loops upon loops [5] (see Fig. 3). This way – we will argue – “meaning” (semantics) can be objectively (i.e., in a measurable way) constructed by the *intrinsic observation process performed by the outer loop(s) on the inner loop(s).* When the agent has come to life there is only the innermost loop working. This loop has a desired state, which, however, can not be maintained all the time (weak homeostasis). These deviations from the desired state of the first loop are observed by the second loop. Thus, the second loop can adjust itself in a way that the deviations from the desired state of the inner loop are minimised. This observation process takes place implicitly and happens strictly *inside* the agent (without any *external* evaluation). From the moment a second loop has been established it can be observed by a third loop and so on. Thus, an agent-centred semantic is developed which has only a meaning for the agent itself.

The nested observation-process becomes more clear with an example of building a secondary loop (Fig. 3) on top of the primary reflex loop in our reactive agent (Fig. 1 a and Fig. 2). The primary loop receives touch- and the secondary loop visual input. When the agent moves around, obstacles will lead first to a vision signal and only later to a touch signal as soon as the agent hits them. 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 [25]). After learning the agent produces an active avoidance movement before touching an obstacle.

How does “meaning” arise 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 this point in time there are signals and actions which are “earlier” and others which are “later”. 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. While this is not explicitly done in such a simple agent, it still represents a very simple intrinsic, semantic evaluation. The outer loop implicitly attributes a meaning to a reaction (e.g. “earlier” or “later”) with respect to the temporal frame of reference given by the inner (reflex) loop.

It is also very important to realize, that there is no external evaluation (by means of a punishment or reward signal) necessary for this type of learning. Learning is strictly correlative and takes place without the interference of an external observer.

Furthermore, we note that this example is certainly not exhaustive and many other ways may exist for achieving the same, but here we have shown that we can still employ a strict inner (organismic) perspective in defining “meaning”.

4 Are robots superior to robot simulations?

So far we have tried to provide evidence that computer programs can be embodied in the sense described above. Furthermore, such structures, when built as nested loops, are able to develop their own internal semantics which was taken as a very first step towards intelligence. The remaining question is: If there is a robot and a robot simulation, both using the same architecture and the same algorithms, will they be indistinguishable or will the properties of the physically embodied robot always be superior to that of the simulation.

One advantage of a technical systems is that we have access to all inner states and can address this question - as before - “from the inside”. What does indistinguishable (or even identical) mean with respect to inner states? Certainly the synaptic weights that develop in the robot and in the simulation will never be identical. Only when using similar initial conditions (and a similar environment) they would be roughly similar, but this does not tell much. Having to adjust the environment is a questionable procedure when wanting to show that robot and simulation are *in essence* identical.

Thus, the question cannot be so easily answered by just looking at structural similarities. One has to go back one step and consider that robot and simulation both develop into the same weak homeostasis with their environment. For the robot-simulation the environment is very simple (and defined in algorithmic terms only), the environment of the real robot is by far more complex. However, both robot and simulation only receive extremely restricted sensorial inputs from their environment and both experience (essentially) the same (collision- and range finder) signals. Thus, the richness of the real environment, which we experience with all our senses, remains hidden from the simulated as well as from the real robot. Nonetheless, both develop into the same homeostasis. This is due to the fact that they have both identical desired states (“no collision”) and a *robustly working reflex* to begin with. Thus, both robot and simulation are identically embodied. Differences could only arise from the structure of the environment with which they are confronted. Above we have shown that the environment is exclusively defined via the disturbances that arrive at the agents. Thus, if we can assure that the fundamental properties of the disturbances are identical for robot and simulation then both would indeed be indistinguishable.

This has interesting practical implications. The robustness of the employed feedback loop structures is due to the fact that only very limited knowledge of the environment is needed to set it up.

> If one is able to create a robot simulation which attains weak homeostasis then one should also be able to successfully transfer this simulation to the real world. This assumes that the fundamental properties of the arriving disturbances are the same.

What are the *fundamental properties* of the disturbances? Here we note, that any disturbance once it has entered the agent turns into a (sensor-)signal. Given that robot and simulation are in our case identically embodied one only needs to analyse the properties of these signals (and not anymore the complex properties of the world) in order to find out if the disturbances are identical, similar of different for robot and simulation. Thus, this viewpoint — the inner perspective — is again helpful for the engineer to assure that a simulation and its corresponding robot are indeed in principle identical.

From the perspective of control-engineering the term *weak homeostasis* can be made more specific. Weak homeostasis means that a system remains stable under the influence of disturbances. Hence, the values of the variables of the system have to remain bounded. Engineering has provided us with powerful tools which allow us to analyse the stability of systems such as those depicted in Fig. 1a [9, 19, 22]. Stability is guaranteed if the disturbance D only leads to bounded values in the loop. For example, one can choose V and ask if V remains bounded if a certain disturbance D enters the loop. In the Laplace domain this leads to:

$$V = DG = D \frac{P_0 H_0}{1 - P_0 H_0} \tag{1}$$

Stability is usually guaranteed if the poles of this equation remain on the left half plane of the transfer function G. This assures that any disturbance causes a damped reaction of the feedback towards the desired state.

5 Observer Problems

Up to this point we have radically employed the perspective of the organism. In this section we want to explore in the same radical way the other perspective, namely the perspective of an external observer. It is an interesting question, how an autonomous agent appears to an external observer who can only observe the *behaviour* of the agent. The *input-processing* and therefore also the feedback-loop is not observable. Can we through such an observation process assess the inner states of an agent? Can we make clear-cut observations if an agent might be autonomous or intelligent? Can we draw conclusion concerning the similarity between robot and robot simulations?

This and the next section we will show that the observer perspective has to be very carefully applied in trying to answer such questions. In particular, one should now *rigorously* apply the observer perspective carefully avoiding to mix inner and outer views which may have led to several fallacies in the treatment of the above problems.

5.1 Can we draw conclusions about inner states?

With some frustration, we find that even in the case of our very simple agent the answer to this question is no, because we have constructed two examples of either an attraction or a repulsion behaviour based on the same desired inner state, namely that of a reflex-avoidance.

In Fig. 4 we show a robot simulation where the robot is learning to avoid obstacles and, at the same time, to find food, simulated by the disks. The learning process for obstacle avoidance is the same as above. The desired inner state of this circuit is to keep the signals at the collision sensors zero; hence to avoid the retraction reflex. Food finding is learned in the following way. The robot has two simulated taste-sensors at its two front corners. At the same locations it also has two sound sensors. We simulate the case that food-sources also produce noise[5] indicated by the circular fields which surround each food-source.

The robot has a built in taste-reflex. Whenever it tastes (i.e. touches) a food source with one of its taste-sensors it will turn in the direction of the body corner where the touched sensor is located, thereby driving more centrally into the food source. Thus, a sign is associated to each of the sensors and we use the difference signal Left-Right to control the turning. No turning reflex will be elicited if the robot moves straight into the food source because then both taste sensors will respond at the same time and the difference signal is zero. The food sources are removed after they have been touched.

The robot learns to produce an anticipatory turning reaction in response to

[5] As a sideline we note that this actually many times happens in reality when a prey moves and rustles in the shrubs.

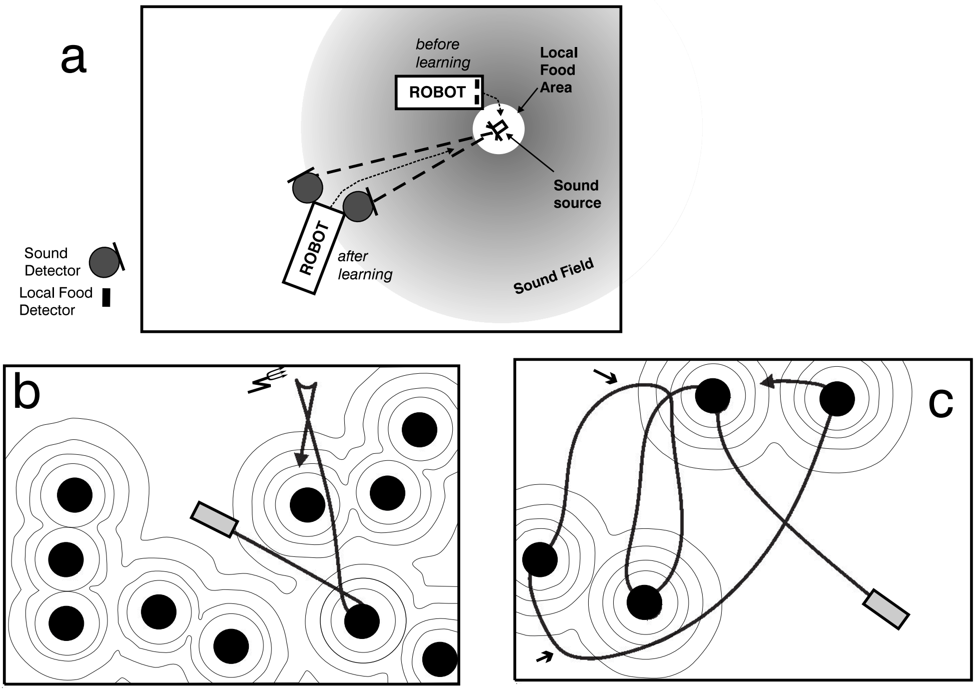

Figure 4: Comparing two paradigms. a) An arena is shown which contains food-sources (disk) surrounded by a sound field. Food sources are not regarded as obstacles and they are moved out of way when touched (i.e., when "eaten") and reinserted somewhere else. Only at the end of the simulation food sources finally disappear (see c). Panel (b) depicts the situation before and panel (c) that after learning. Arrows indicate successful (plain arrow) or unsuccessful (devil's arrow) obstacle avoidance, which for an observer looks like a "punishment-avoidance" paradigm. Whenever the robot has driven through a disk it has obtained "food". This could be interpreted as a "reward-retrieval" paradigm. In (b) both actions occur as the consequence of two reflexes. Obstacle avoidance by means of a retraction reflex (negative motor signal) and food retrieval by an attraction reflex (positive motor signal). In (c) both reflexes are replaced by anticipatory actions. Thus, in both cases the inner state is to avoid either reflex. Punishment and reward associations are not applicable. The complete simulation can be obtained from: `http://www.cn.stir.ac.uk/predictor/animat/`

the sound field. In the same way as the taste-reflex is performed it learns now to turn towards the more strongly stimulated sound sensor. Also here we use a difference signal between the left and the right sound sensor. As a result the robot will finally already from some distance drive straight into the food source. An external observer will find that the robot now performs an anticipatory attraction reaction towards the food, while the same robot is being repulsed by obstacles and walls.

However, what is the desired state of the taste-reflex circuit? Without any difference to the collision-reflex circuit, also here the desired state is to keep the control signal equal to zero thereby avoiding the taste reflex.

Thus, attraction and repulsion behaviour are both learned by avoiding their associated reflex actions. There is no way to draw this conclusion by just looking at the simulations. The inner states of even such a simple agent remain inaccessible by observation.

An important follow-up problem which arises from these observation is that one has to be exceedingly careful in applying the concepts of “punishment” or “reward” to data obtained from behavioural observations. Our examples are both cases of “reflex-avoidance”: Does the one process now avoid a “punishment” while the other retrieves a “reward”? We do not think that such associations can be rigorously applied, neither in such a simple robot, nor in much more complex animals [32]. There is always a (possibly strong) interpretational aspect included in such conjectures.

5.2 Can we draw conclusions about intelligence or autonomy?

From a restrictive point of view the answer to the second part of this question is yes. This, however, requires to accept and not question one certain definition (or dominant sub-aspect) of autonomy: Several authors claim that unpredictability of behaviour is one of the basic properties which shows that an agent is autonomous [41, 1]. Note, this definition is *rooted* in the behavioural domain and therefore it can also be rigorously answered by looking (only) at behaviour.

In our case one can see that during learning the complexity of the observable behaviour clearly grows. The reflex based agent performs very stereotyped reactions when hitting an obstacle. These reactions can be exactly predicted by the observer. After learning the agent has developed a proactive behaviour and his actions have now become largely unpredictable from the outside. The level of behavioural complexity (i.e., unpredictability) which our agents develop is still rather low. Nevertheless, this also shows that such embodied nested loop architectures may represent one possible early step towards autonomous agents which are not bound to a physical substrate. Other more complex aspects of autonomy are not necessarily covered by such an architecture, but we do not see any fundamental reasons why more complex aspects could not be built on top of this.

Concerning intelligence, the situation is similar. Strictly adopting the observer perspective means that one has to accept the significance of a Turing Test. We will not enter the ongoing discussion if such tests are conclusive or not (essentially questioned by Searle [33]). Instead we state that from the perspective of an observer, our agent would have to be judged as intelligent as any other agent whose behaviour is indistinguishable from it. From this it is obvious that the observable "intelligence" of such agents is still fairly low. The development of a simple semantic evaluative inner structure can also only be seen as perhaps one of the first steps towards true intelligence. Important to remember at this stage however is that an external behaviour will not be able to use his/her own semantics in the same way as the agent does. For an observer a certain type of behaviour may "look" better (or worse) than the behaviour prior to learning. These same attributions may not at all apply to the agent, though. This problem is similar to the problem of punishment-reward associations discussed above.

5.3 Can we distinguish robots and robot simulations?

From the perspective of an external observer robot and simulation will have to be regarded as identical if, after in-depth observation, there are no *behavioural* differences to be discovered. This amounts again to a Turing Test.

Fig. 5 shows four movement traces obtained with a robot and a simulation in the obstacle avoidance task. Two traces show the animats before and after anticipatory learning. The algorithm used was fully identical in simulation and the real robot also using the same parameters for the initial setup and a similar environment. However, even such a simple task (i.e. obstacle avoidance) leads always to different solutions for robot and simulation. In the shown examples this is very obvious because different starting positions were chosen, but, even if those were the same one would observe that even a small initial difference in the setups will lead to growing differences in the runs. This, however, is not the point. The question is rather, are there fundamental, observable differences between robot and simulation. Here, we challenge the reader to tell us which set of traces comes from the simulation and which from the real robot[6]. We think these traces do not show any fundamental difference. Of course this example is exceedingly simple,

[6]Solution: a) and c) are the traces from the a real robot and b) and d) are the traces from a robot simulation.

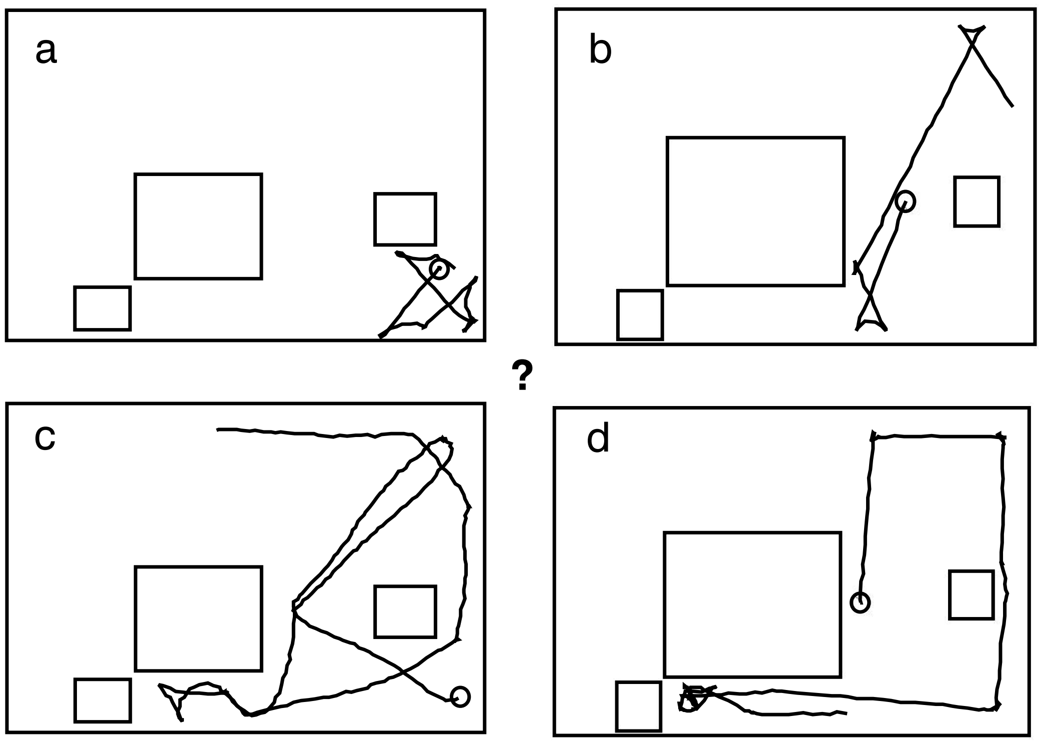

Figure 5: Can we distinguish between the simulation and the real robot? What is the simulation and what is the real robot experiment? Part A to D show the traces of a simulated robot and a real robot before and after learning. The traces were obtained by hand for both experiments by clicking on the robot-image. Learning was performed by the ISO-learning algorithm [25] which turns the avoidance-reflex into a pro-active avoidance action.

but we would like to direct the reader's attention to the existing state-of-the-art animations of the movie industry. These artificial dinosaurs look and behave so realistically that behavioural observation does not reveal that they are just simulations.

When considering animals or other humans, behavioural observation is the only way to address these questions because we do not have access to their inner states. As a consequence it has long been known in field of psychology that it is exceedingly hard to "correctly" interpret behavioural studies [28]. Simulated behaviour in form of robots or its equivalent simulations can overcome this problem as all their internal states are accessible and therefore the sensor-motor loop can really be closed and interpreted. Thus, robotics can help to overcome observer-problems.

6 The failure of traditional AI

So the question why AI has encountered so many problems should be re-addressed here. We believe that it is not the failure of algorithms to be embodied in any (physical) sense which is the underlying reason. Instead, we think that a central problem lies in the traditional AI-approach as such. Conventional AI has always adopted an engineering (observer) perspective. Algorithms were treated as input-output systems and only the engineer received feedback about the success or failure of the algorithm. In a simulated AI input/output system a large but limited number of input conditions exists and an external observer will still be able most of the time to hand-adjust the algorithm to produce an I/O-reaction which he/she deems correct. This way the engineer is building a forward-model of the simulated world into the system, which is possible as long as the complexity remains limited and foreseeable.

Transferring the successful simulation to the real world, however, now introduced unforeseeable contingencies and many times the algorithms failed. This was taken as evidence that only physically embodied agents which can be rooted in the environment will be able to succeed and perform "correctly" (or "appropriately"). We believe that this was the wrong conclusion. Instead this article has argued that it is indeed the lack of feedback which leads to the failure of engineered I/O-systems. The algorithm does not receive any information about its own success and hence cannot adapt its own reactions. As a first order problem this is due to the fact that the engineer, who *does* receive the feedback, cannot in the real-world decide which aspects of the complex sensor-motor action cycles has led to the wrong results. The complexity of the real world prevents designing complete forward models.

There is, however, a second problem with the I/O-perspective. In general one finds that engineers and agents may have different intentions which relates to different desired states. We call it the *Second Chicken & Egg Problem*: The farmer (engineer) wants to maximise the egg-production output and hence removes the freshly laid eggs. The chicken wants to maximise the sit-on-the-egg-time in order to facilitate breeding.

Let us have a look at the chicken-egg problem #2 in a more formal way (see Fig. 6), namely as a cybernetical feedback system [15]. Let us first describe it in the classical way and then identify which variables and transfer functions are

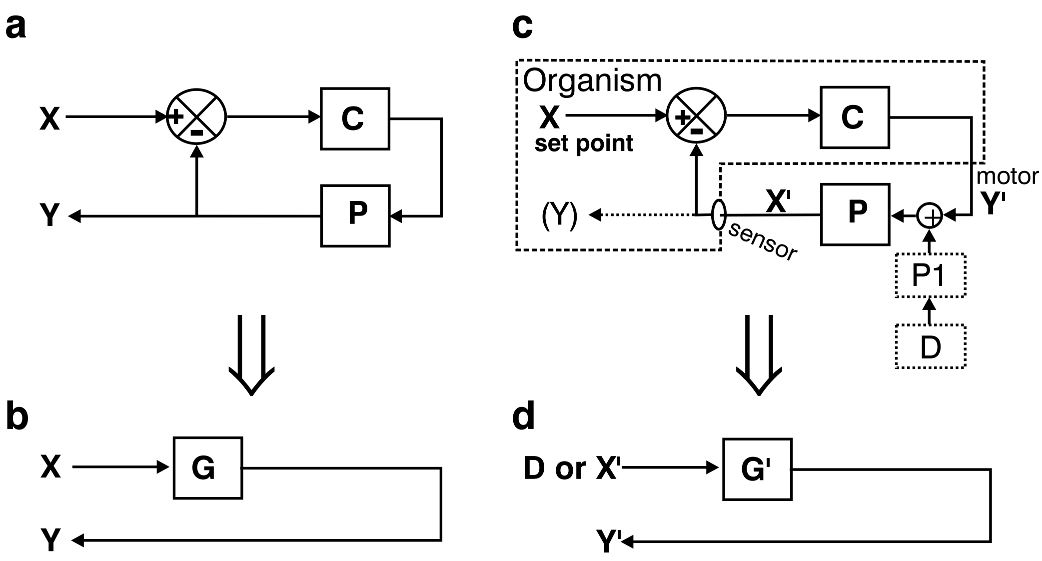

Figure 6: The chicken egg-problem #2 in cybernetical terms. a) A typical cybernetical control setup. The input X is filtered by the feedback-system and generates an output Y. Note that all components are observable (see [12]). b) The feedback system can be collapsed into a single transfer function: $G = \frac{CP}{1+CP}$ c) A cybernetical control setup of an organism. The organism wants its input X' to take the value defined internally by the set point X. This should be achieved by the motor output Y'. Only Y' is observable from the outside. From the organism's point of view $X, X' and Y$ are observable. d) An external observer can also try to collapse the organism into an I/O system, for example by performing stimulus-response experiments and trying to deduce G'.

observable from the organisms point of view (the chicken) and from the external observer, namely from the farmer or engineer [16, pp.133–135].

In classical feedback control all variables are observable per definition [12, 2]. The feedback loop in Fig. 6a represents such a feedback control system. The transfer function C of the controller is also known and the transfer function of the plant P can be at least approximated with appropriate test signals. The input X of the system is given by the engineer in form of a set-point. The output Y should follow the input X in an optimal way. As a consequence it is in principle possible to collapse the feedback-loop into an input-output transfer function $G = \frac{CP}{1+CP}$. Even if the transfer function P is not given, the system can still be physically designed as a feedback loop — as drawn above — because such a system does not need precise knowledge of P. This is the advantage of feedback control in general and we have mentioned it before.

Let us now look at the system shown in Fig. 6c which represents the situation we are facing when we are dealing with an organism. We have to decide if we are employing the perspective of the organism or the perspective of the external observer. The perspective decides which variables are observable and which are not. We find that from the organism's point of view the variables X and $Y = X'$ are observable. Here we have for simplicity neglected any possible signal transform at the sensor which renders $Y = X'$. This is basically the same situation as that shown in Fig. 6a with the exception that Y' is not directly observable by the organism because this signal will only arrive at its sensors after having been filtered by the environment P. Again we note that knowledge of this signal is not necessary for the function of the feedback loop. In terms of the chicken Y' is the egg output. $X' = Y$ is the sensation of the egg and X is the set-point which defines that sensing an egg is the desired state (in contrast to not sensing an egg). The transfer-function G defined by the chicken itself is optimal if in case of a disturbance D (the farmer) the loop instantly reestablishes the desired state, namely by producing a new egg. The same applies if the chicken is initially in the state of not wanting an egg (set-point is zero) and then at some point decides that "it would be nice to have an egg now" (changing the set-point to one). If the loop establishes quickly the desired state then it operates in an optimal way.

For the farmer/engineer there are different variables observable. He/she can only observe the output of the chicken Y'. All the other variables are hidden in the organism and cannot be observed. However, this is sufficient for the farmer as he/she is only interested in the *output* of the chicken. Note that there is no way of easily defining the transfer function G' of the whole system from this outside perspective (Fig. 6d). One has to guess the input-channels which are part of the feedback-loop, one has to guess the environmental feedback, the disturbance and the set-point. Only with this knowledge one could try to define a transfer function which predicts how the organism reacts to a disturbance D: $Y' = \frac{CX - CP_1D}{1 - CP}$ or to an input X'. Note, that the latter paradigm: Trying to assess G' by measuring the output Y' in response to an input X' represents the traditional and widely used approach of treating an organisms (or a sub-system) as a stimulus-response system. Most physiological and many psychological experimental paradigms follow this approach, which may be problematic as it ignores the feedback, which — at least in a non-laboratory situation — plays a major role for the organism.

Thus, the intentions of an external observer may be quite different from that of an agent. While the agent tries to keep its *input* at a certain level the observer desires to have the agent's *output* at a certain level.

The chicken-egg problem #2 can be summarised in the statement that:

> Autonomous systems control their inputs and not their outputs [40].

This is why deep-inside true engineers should never like autonomous systems which hide their variables and transfer functions.

As a consequence the engineer outside will not only build an *incomplete* forward model, even worse, he/she might have built from the viewpoint of the agent the *wrong* forward model of the world. The right forward model can only be created by the agent itself as the agent evaluates at its *inputs* as to whether the forward model has been successful or not.

Learning plays a crucial role in developing the optimal forward model *for the agent* as it can develop the forward-model from its own perspective. It should be emphasised that also learning is guided only by inputs, namely by deviations from the desired state measured at the agent's sensors. Therefore, in order to generate true (observer-free or teacher-free) autonomy, only learning rules can be used which are guided by an input-condition and not by an output-condition. Learning rules which are suitable for creating autonomy belong to the class of unsupervised, Hebbian (i.e., correlation based) learning-rules. Note, that on the other hand most of the existing learning rules are guided by an output condition, especially all reinforcement learning rules like TD-learning [37] and all supervised rules like the delta-rule [43]. Very complex goals can indeed be achieved with these rules, but those are always the goals of the network designer and normally not the goals of the agent.

7 Practical implications

We have tried to make clear that a very broad concept of embodiment may suffice as a necessary condition for intelligence. Furthermore, we have shown how such systems should operate in conjunction with their environment. Essentially, to use a more novel term, how they should be situated (or rooted) in their environment. This has been done by applying concepts from control- and systems-theory.

We believe that there are two aspects in this study which could be of a somewhat longer lasting relevance.

1. As opposed to other approaches, a system theoretical approach provides the opportunity of mathematical treatment in three respects:

(a) Creating an agent means employing *stable feedback loops* in order to establish desired states in a weak homeostasis [22]. This implies that (non-lethal) deviations from the desired states are permissible and even *wanted* in order to keep the system open to its environment.

(b) Forward-models [20] of the above feedback-loops have to be created by the agent *itself* in order to improve homeostasis. That this is possible has been successfully demonstrated by rigorously proving that temporal sequence learning (using the ISO-learning algorithm) is able to generate a forward-model of a reflex [24].

(c) The degree of embodiment can be measured in this context by the performance of the boundary. Possible measures are Ashby's *requisite variety* [2] or Polani's relevant information [23] which are both suitable for a closed-loop context.

2. Furthermore, through such a treatment the problems of applying the inner- versus the outer perspective can be more clearly discussed. Living agents do not reveal their desired states and their strategies how they control their inputs. However, their output is observable and tempt the observer to treat a living being (or an autonomous robot) as an I/O-system. Robotics, however, gives us the opportunity to observe even closed-loop systems because all the components (inputs, desired states, etc.) of such loops can be uniquely associated to observable (physical) structures and/or processes.

Having said this, it is quite clear to us that the systems discussed in this study are still terribly simplistic. It will be a matter of future research to show if the concepts presented here will survive under more complex conditions.

Acknowledgements

We are grateful to Leslie Smith, Barbara Webb, Peter Sutherland and to the members of the CCCN-seminar for their helpful comments during various stages of this work. This study was supported by grants from SHEFC RDG "INCITE".

References

[1] David Anderson. *Artificial Intelligence and intelligent systems: The implications*. Ellis Horwood LTD, Chichester, England, 1989.

[2] W. Ross Ashby. *An introcduction to cybernetics*. Methnen+Co LTD, London, 1956.

[3] Ludwig von Bertalanffy. *General Systems Theory*. Braziller, New-York, 1973.

[4] Rodney A Brooks. How to build complete creatures rather than isolated cognitive simulators. In K. VanLehn, editor, *Architectures for Intelligence*, pages 225–239. Erlbaum, Hillsdale, NJ, 1989.

[5] Rodney A Brooks. A robot that walks; emergent behaviors from a carefully evolved network. Technical Report 1091, MIT AI Lab, 1989.

[6] Rodney A. Brooks. Intelligence without representation. In Haugeland John, editor, *Mind Design II*, chapter 15, pages 395–420. MIT-press, Cambridge, Mass., 1997.

[7] Rodney A. Brooks. The relationship between matter and life. *Nature*, 409:pp. 409–411, 2001.

[8] K. Dautenhahn, B. Ogden, and T. Quick. From embodied to socially embedded agents - implications for interaction-aware robots. *Cogn. Syst. Res.*, 3:397–428, 2002.

[9] John J. D'Azzo. *Linear Control System analysis and design*. Mc Graw, New York, 1988.

[10] Daniel C. Dennett. Cognitive wheels: The frame problem of ai. In C. Hookway, editor, *Minds, machines, and evolution*, pages 129–151. Cambridge University Press, Cambridge, 1984.

[11] S Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.

[12] R. E. Kalman. Mathematical description of linear dynamical systems. *j-J-SOC-IND-APPL-MATH-A-ON-CONTROL*, 1:152–192, 1963.

[13] Niklas Luhmann. *Social Systems*. Stanford University Press, Stanford, California, 1995.

[14] Humberto Maturana and Francisco J. Varela. *Autopoiesis and cognition: the realization of the living*. Reidel, Dordrecht, 1980.

[15] David John McFarland. *Feedback mechanisms in animal behaviour*. Academic Press, London, 1971.

[16] David John McFarland. *Problems of animal behaviour*. Longman Scientific & Technical, Harlow, 1989.

[17] David John McFarland. *Intelligent behavior in animals and robots*. MIT Press, Cambridge, MA, 1993.

[18] Clare D. McGillem and George R. Cooper. *Continous and discrete signal and system analysis*. CBS publishing, New York, 1984.

[19] Norman S. Nise. *Control Systems Engineering*. Cummings, New York, 1992.

[20] William J. Palm. *Modeling, Analysis and Control of Dynamic Systems*. Wiley, New York, 2000.

[21] Rolf Pfeifer and Christian Scheier. *Understanding Intelligence*. MIT Press, Cambridge, MA, 1999.

[22] Charles L. Phillips. *Feedback control systems*. Prentice-Hall International (UK), London, 2000.

[23] Daniel Polani, J. T. Kim, and T. Martinetz. An information-theoretic approach for the quantification of relevance. In J. Kelemen and P. Sosik, editors, *Advances in Artificial Life*. Springer, 2001.

[24] Bernd Porr, Christian von Ferber, and Florentin Wörgötter. Iso-learning approximates a solution to the inverse-controller problem in an unsupervised behavioural paradigm. in press: Neural Comp., 2003.

[25] Bernd Porr and Florentin Wörgötter. Isotropic sequence order learning. in press: Neural Comp., 2003.

[26] Ilya Prigogine. *Vom Sein zum Werden*. Piper, München, 1982.

[27] Tom Quick and Kerstin Dautenhahn. Making embodiment measurable. In *Proceesings of the 4th Fachtagung der Gesellschaft für Kognitionswissenschaft*, Bielefeld, Germany, 1999.

[28] Howard Rachlin. *Behaviour and Learning*. Freeman and Company, San Francisco, 1976.

[29] A. Riegler. When is a cognitive system embodied? *Cogn. Syst. Res.*, 3:339–348, 2002.

[30] Robert Rosen. *Life itself : a comprehensive inquiry into the nature, origin, and fabrication of life*. Columbia University Press, New York, Oxford, 1991.

[31] Oliver Sacks. *The Man Who Mistook His Wife For a Hat*. Simon & Schuster, New York, 1998.

[32] Wolfram Schultz, Peter Dayan, and P. Read Montague. A neural substrate of prediction and reward. *Science*, 275:1593–1599, 1997.

[33] John R Searle. Minds, brains and programs. *The Behavioral and Brain Sciences,* 3:417–424, 1980.

[34] Noel E. Sharkey and Tom Ziemke. The new wave in robot learnin. *Robotics and Autonomous Systems*, 22(3-4), 1997.

[35] G. Spencer Brown. *Laws of Form*. George Allen und Unwin, 1969.

[36] John L. Stewart. *Fundamentals of signal theory*. Mc Graw-Hill, New York, 1960.

[37] R.S. Sutton. Learning to predict by method of temporal differences. *Machine learning*, 3(1):9–44, 1988.

[38] Heinz von Foerster. On self-organizing systems and their environments. In M.C. Yovits and S. Cameron, editors, *Self-Organizing Systems*, pages 31–50. Pergamon Press, London, 1960.

[39] Heinz von Foerster. *Sicht und Einsicht: Versuche zu einer operativen Erkenntnistheorie*. Vieweg, Braunschweig, 1985.

[40] Ernst von Glasersfeld. Learning and adaptation in constructivism. In Leslie Smith, editor, *Critical Readings on Piaget*, pages 22–27. Routledge, London and New York, 1996.

[41] W. Grey Walter. *The Living Brain*. G. Duckworth, London, 1953.

[42] Juyang Weng. A theory for mentally developing robots. In *2nd International Conference on Development and Learning*, Cambridge, MA, 2002. IEEE Computer Society.

[43] G Widrow and M.E. Hoff. Adaptive switching circuits. *IRE WESCON Convention Record*, 4:96–104, 1960.

[44] Norbert Wiener. *Cybernetics — or control and communication in the animal and the machine*. The M.I.T. Press, Cambridge, Massachusetts, 2 edition, 1961.

[45] Tom Ziemke. Are robots embodied? In *First international workshop on epigenetic robotics Modeling Cognitive Development in Robotic Systems*, volume 85, Lund, 2001.