

“Value Signals” and Adaptation: An Exploration in Evolutionary Robotics

Marieke Rohde and Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR)
Department of Informatics, University of Sussex, Brighton, BN1 9QG, UK
{m.rohde,ezequiel}@sussex.ac.uk

Abstract. Pfeifer and Scheier write: “If the agent is to be autonomous and situated, it has to have a means of ‘judging’ what is good for it and what is not. Such a means is provided by an agent’s value system” ([8], p. 315). What can it mean for a system to generate “values”? In this paper, we take a closer look at this question. A series of minimal evolutionary robotics experiments, in which an agent is evolved to generate a signal that corresponds to its level of performance, in analogy to the idea of a value system, is presented and discussed, pointing out the essential role of sensorimotor coupling for the integrated process of judgment. The emphasis of the discussion is on the relation between function and mechanism and aims at questioning our intuitions about value systems and the neural correlates of meaningful events and processes.

1 Introduction

What is the relation between function and mechanism? Research in evolutionary robotics is frequently motivated by a scepticism about an isomorphic relation between the structure of behaviour and the structure of the physical mechanism that brings it about (e.g. [11]).

This paper follows this tradition of questioning the localisation of a function. We critically investigate what is called a “value system” and “value guided learning”, mooting the assertion that an encapsulated system can be held responsible for meaningful judgment. Our method allows us to generate a minimal embodied system where “value judgements” result from a coupling between mechanism and behavioural dynamics. We show how the de-compositional view misses out on crucial aspects of how the system works. This example illustrates an important theoretical possibility that traditional approaches to the question of value are unable to account for.

2 Value System Architectures

2.1 What Are Value System Architectures?

What we refer to as “value system architectures” is a class of models for life-time adaptation, characterised by a functional and structural division between

behaviour-generating mechanisms and mechanisms of adaptation. In particular, these models feature a *value system*, which generates a bipolar performance signal directing adaptive processes (value guided learning). Their activity can be seen as the internal generation of a reinforcement signal.

The label “value system” has been taken from the theory of neuronal group selection (TNGS) by Edelman et al. (e.g. [3]). TNGS proposes ontogenetic Darwinian-style evolution as principle of neural organisation ([3], p. 242). A value signal, generated by a value system, is the criterion to reinforce successful behaviour by strengthening the participating synaptic connections, a process akin to natural selection. For instance, a value system for reaching would become active (“good”) if the hand comes close to the target [10].

This underspecification permits more behavioural flexibility than pre-specified motor programmes and allows an organism to manage the effects of anatomical variations on neural control. However, it requires the behaviour generating mechanisms themselves to be value-agnostic and blindly obey the value system’s judgment. The value system is a separate system, which in itself is not supposed to adapt, at least not through its own judgment¹. Value systems are thought to be “already specified during embryogenesis as the result of evolutionary selection upon the phenotype” ([10], p. 968).

This idea of value guided learning has also been transferred to autonomous robotics. Pfeifer and Scheier refer to TNGS in their book “Understanding Intelligence” and support the claim that self-supervision through value systems is essential to direct processes of self-organisation in autonomous agents ([8], p. 467, see Verschure et al. [12] for an example application).

2.2 What Is the Problem With Value System Architectures?

Our argument can be seen as a special case of an argument that others have made before us: It is the claim that the dynamics of behaviour and the dynamics of behavioural learning, even though they can be functionally distinguished and occur on different time scales, need not be brought about by different physiological structures. Simulated evolutionary robotics experiments, first by Yamauchi and Beer [13], then by Tuci, Quinn and Harvey [11], have helped to illustrate this point, by demonstrating how a unitary fixed weight control network can realise fast changing motor behaviour as well as long term modulation of this behaviour (learning).

These existence proofs, even though they teach us to be careful not to presuppose a functional modularity, do not exclude the empirical possibility of such structures. The developmental psychologist Julie Rutkowska [9], however, provides more practical reasons to be sceptical of value system architectures. She argues that “[increased] flexibility requires some more general purpose style of value” ([9], p. 292) than a value module could provide, even though such circuits may work in specific cases. She laments their vulnerability and their restrictive

¹ Some authors hold it possible “that different value systems interact, or that hierarchies of specificity might exist.” ([10], p. 969).

semantics consequent to the built-in evaluation criteria. A similar limitation is pointed out by Pfeifer and Scheier, who describe a “trade-off between specificity and generality of value systems” ([8], p. 473): A very specific value system will not lead to a high degree of flexibility in behaviour, while a very general value system will not constrain the behavioural possibilities of the agent sufficiently.

The common denominator of these different issues raised by different researchers is summarised in Rutkowska’s question of whether a value system constitutes a “vestigial ghost in the machine” ([9], p. 292). A value system that applies pre-specified evaluation criteria to pre-specified sensory states to steer ontogenesis in a top-down manner, even if it guides the adaptation of real-time situated and embodied behaviour, is in itself a disembodied control structure. As such, it suffers from all the problems associated with traditional disembodied artificial intelligence architectures, which have been pointed out many times (e.g. [2, 7, 8]): They are rigid and non-adaptive, their functionality relies on the intact functionality of dedicated input and output channels and they can only deal with scenarios that could be foreseen when they were designed.

2.3 The Only Good Ghost Is a Dead Ghost

The astonishing fact about value system architectures is that, despite the outlined disembodied nature of the value system, these architectures are very popular with researchers that share our concerns about situatedness and embodiment in the study of intelligent behaviour, and who are deeply sceptical towards classical symbolic approaches. For instance, Sporns and Edelman point out how TNGS models, through their increased flexibility, can overcome difficulties such as anatomical variations, which are “challenging to traditional computational approaches” ([10] p. 960). It is probably unquestioned that “Understanding Intelligence” by Pfeifer and Scheier [8], the very volume that advertises value guided learning, is one of the most important books to promote the situated and embodied approach.

Maybe, it is “shrinking” the homunculus that makes the difference for these researchers, after all, value systems are just a *vestigial* ghost in the machine². Maybe, empirically, there are “simple criteria of saliency and adaptiveness” ([10], p. 969) that can *a priori* specify what will be good and what will be bad *a posteriori*³.

As a neuroscientific theory, TNGS is backed with empirical evidence. There is e.g. a correspondence between salient events in the environment and the activity of cell assemblies in the brain stem and the limbic system that modulate synaptic

² For instance, Edelman’s statement that “[TNGS] relies only minimally upon codes” ([4], p.45) suggests this interpretation.

³ An option that we can probably exclude is that “value” and “value systems” are simply ambiguous terms and used to describe phenomena on both the mechanical and the functional level. When Edelman maintains e.g. that “general information about the kinds of stimuli that will be significant to the system is built in” ([3], p.58), it is obvious that a literal reduction of function to mechanism underlies the idea of value system architectures.

changes in the cortex[5]. The bigger question to be asked in this context is: What can we deduce from such a correspondence⁴?

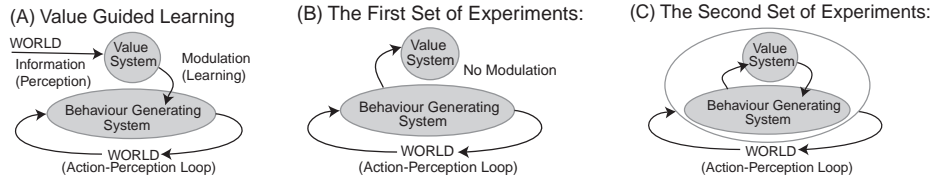


Fig. 1. Schematic view of value system architectures (A), and the alternative views resulting from our first (B) and second (C) set of experiments.

The simulation experiments we present in this paper approach this question in minimal controlled settings. A mobile agent is designed through artificial evolution to perform simple phototaxis and, at the same time, to generate a signal that corresponds to its level of performance. There is no *a priori* need or function associated with this estimate, it simply serves us as analogy with the aforementioned brain structures identified as value systems. In a first set of experiments, a value signal is generated that has no effect on the network dynamics. With this experiment, we raise the question whether it is adequate to think of value generation as the application of a pre-specified function, which can be separated from sensorimotor behaviour (as in Fig. 1 (A)), or if judgment is rather an activity just like phototaxis and is constituted within a closed sensorimotor loop (Fig. 1 (B)). In a second set of experiments, the internally generated value signal is fed back into the neural dynamics of the agent (Fig. 1 (C)). With this experiments we want to question intuitions about the value system *modulating* the behaviour dynamics. We emphasise the consequences of the reciprocal causal links that go in both directions, not only top-down from the value system to the behaviour generating network.

3 The Model

The model is deliberately minimalist. It does not aim to model actual brain structures, as the cited models, it serves to illustrate a conceptual argument.

A circular two-wheeled agent of 4 units diameter is designed by evolutionary search to perform phototaxis. The control networks evolved are continuous time recurrent neural networks (CTRNNs, e.g.[1]) with variable size and structure (see

⁴ According to Kandel, the “idea that different [brain] regions are specialized for different functions is now accepted as one of the cornerstones of modern brain sciences” ([6], p. 9). We think that such functional specialisation of brain regions is questionable, at least as a general case.

below). The dynamics of neurons n_i in a CTRNN of N neurons are governed by

$$\tau_i \frac{da_i(t)}{dt} = -a_i(t) + \sum_{j=0}^N c_{ij} w_{ij} \sigma(a_j(t) + b_j) + I_i \quad (1)$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the standard sigmoidal function and I_i is the external input to n_i . The weights $w_{ij} \in [-8, 8]$ from n_j to n_i , the bias $b_i \in [-3, 3]$ and the time constant $\tau_i \in [16, 516]$ are determined by a genetic algorithm (GA). C is the $n \times n$ connectivity matrix with $c_{ij} = 1$ if there is a connection from n_j to n_i and $c_{ij} = 0$ otherwise.

The agent has two sensors $S_{L,R}$ with an angle of acceptance of 180° , which are oriented towards $+60^\circ$ and -60° , with added uniform directional noise $\epsilon \in [-2.5^\circ, 2.5^\circ]$. Their activation is fed into input neurons by $I_{S_i}(t) = Sg \cdot S_{L,R}(t)$ with Sg evolved $\epsilon [0.1, 50]$ and $S_{L,R}(t) = 1$ if the light is within the sensory range of $S_{L,R}$ at time t and $S_{L,R}(t) = 0$ otherwise. Note that the binary character of the light activation makes the estimation of the distance to the light non-trivial. The motor velocities are set instantaneously at any time t by $M_{L,R}(t) = M_G \cdot (z_{M_{i+}}(t) - z_{M_{i-}}(t)) + \varepsilon$ where M_G is the motor gain $\epsilon [0.1, 50]$. $z_{M_{i\pm}}(t)$ is the neural output of one of the two neurons controlling $M_{L,R}$ and $\varepsilon \in [0, 0.2]$ is uniform noise. A fifth output neuron generates the performance estimate $E(t) = z_{M_5}(t)$.

The connectivity C and the size of the network is partially evolved. Connections to input neurons or from output neurons are not permitted. Input neurons can project to output neurons and to hidden neurons, hidden neurons can project to other hidden neurons and to output neurons. The network can have varying numbers (0–5) of hidden neurons. In experiments where the value signal E is integrated into the network dynamics (Sect. 4.2), the estimator neuron changes status to become another interneuron. In some experiments, parts of the network structure and parameters were excluded from continued evolution at a certain stage.

Parameters for the control network are evolved in a population of 30 individuals with a generational genetic algorithm with real-valued genes $\epsilon [0, 1]$, truncation selection ($\frac{1}{3}$), vector mutation [1] of magnitude $r = 0.7$ and reflection at the gene boundaries. The sensor gain S_G , the motor gain M_G and the time constants τ_i are mapped exponentially to the target range. The existence or non-existence of hidden neurons and neuronal connections is determined by the step functions $x > 0.7$ and $x > 0.6$ respectively. All other values are mapped linearly to their target range.

In every evaluation, the robot is presented with a sequence of 4-6 light sources that are placed at a random angle and distance $\epsilon [40; 120]$ from the robot. Evaluation trials last $T \in [3000, 4000]$ time steps. They are preceded by $T' \in [20, 120]$ simulation time steps without light or fitness evaluation, to prevent that the initial building up of activity in the estimator neuron follows a standardised performance curve. Each light is presented for $t_i \in [\frac{T}{5} - 100, \frac{T}{5} + 500]$ time steps. The network and the environment are simulated using the forward Euler method with a time-step of 1 time unit.

The fitness $F(i)$ of an individual i is given by

$$F(i) = F_D(i) \cdot F_E(i) + \varepsilon F_D(i) \quad (2)$$

where $F_D(i)$ rates the phototactic behaviour and $F_E(i)$ rates the fitness prediction. The second term ($\varepsilon = 0.001$) is included to bootstrap the evolution of behaviour, as the coevolution of light seeking and estimation of performance from scratch is difficult for evolutionary search. $F_D(i)$ is given by

$$F_D(i) = \frac{1 - M^2}{T} \int_0^T \max\left(0, 1 - \frac{d(t)}{d(t_0)}\right) dt \quad (3)$$

with $M = \frac{0.125}{T} \int_0^T \frac{M_L(t) - M_R(t)}{M_G}$. $d(t)$ is the distance between robot and light at time t and t_0 the last displacement of the light source. The estimate fitness F_E has gone through a long but necessary process of refinement and complication. It is given by

$$F_E(i) = \sqrt{\max\left(0, \frac{e(\bar{d}, d) - e(E, d)}{e(\bar{d}, d)}\right) \cdot \max\left(0, \frac{e(0, \dot{d}) - e(\dot{E}, \dot{d})}{e(0, \dot{d})}\right)} \quad (4)$$

with $e(x, y)$ the sum of squared error $e(x, y) = \int_0^T (x(t) - y(t))^2 dt$. \bar{d} is the average of $d(t)$ during each trial. $\dot{d}(t)$ and $\dot{E}(t)$ are the derivatives of $d(t)$ and $E(t)$ averaged over a sliding time window $w = 250$ time steps (interval borders for $e(x, y)$ have to be adjusted accordingly).

The evaluation of a network i on $n = 6$ trials is given by

$$F(i) = \sum_{j=1}^n F_j(i) \cdot 2^{-(j-1)} \cdot \frac{1}{\sum_{j=1}^n 2^{-(j-1)}} \quad (5)$$

where $F_j(i)$ gives the fitness on the j^{th} worst evaluation trial for individual i , which gives more weight to worse trials and thereby rewards the generalisation capacity of the evolved networks.

4 Results

4.1 Generating a Value Signal

In this section, we describe and analyse an individual evolved agent. It was selected because of its simplicity and because its way of estimating performance is representative for the most frequently evolved strategy.

The network evolved (Fig 2, (A)) does not have hidden neurons, recurrent connections or slow time constants, i.e. its behaviour hardly relies on internal state and its complexity is minimal, even within the already restricted range of possibilities. For rhetorical reasons, we start with the description of the value system, before we describe the light seeking behaviour.

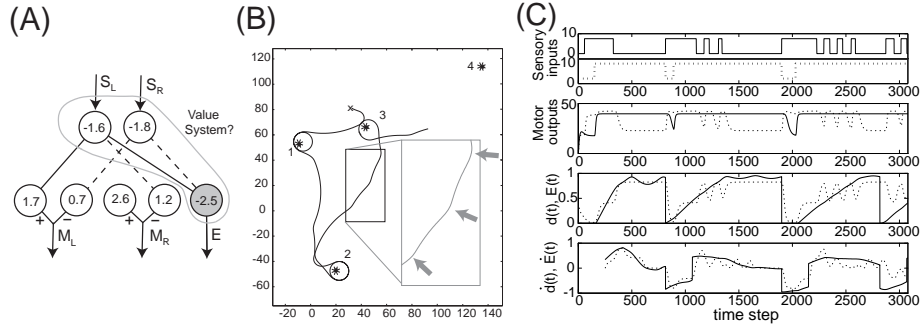


Fig. 2. (A) The distance estimator network (θ in neurons, dotted lines inhibition, solid lines excitation). (B) Trajectory following four presentations of light sources. Arrows indicate the punctuated turns during $t = 2200 - 2700$ (see text). (C) The evolution of different variables over time in the same trial (Top to bottom: $S_{L,R}$, $M_{L,R}$, $d(t)$ vs. $E(t)$, $\dot{d}(t)$ vs. $\dot{E}(t)$).

The neural structures participating in the generation of the value signal are just the two input neurons and the estimator neuron, so if anything, we would have to call this sub-system the value system. In the absence of light, or if the network receives input only on its right light sensor ($S_R = 1, S_L = 0$), it estimates $E \approx 0$. If light is perceived with both sensors, it estimates $E \approx 0.5$, and if the network receives input only in its left light sensor ($S_R = 0, S_L = 1$), the estimate reaches its maximum of $E \approx 0.8$. The judgment criteria of this value system can thus be described as “seeing on the right eye is good, seeing on the left eye or not at all is bad”. Intuitively, these rules do not make sense. Nevertheless, as we can see in Fig. 2 (C) (bottom two plots), both $E(t)$ and $\dot{E}(t)$ (dotted lines) follow with amazing accuracy the actual values $d(t)$ and $\dot{d}(t)$ (solid lines), particularly if we remember the poor sensory endowment of the agent.

The agent’s light seeking behaviour is realised by the network minus the estimator neuron. In the absence of sensory stimulation, the agent slowly drives forward, slightly turning to the right. If $S_R = 1$ and $S_L = 0$, the “brake” on the left motor M_L is released, which leads to a sharper turn to the right. If $S_R = 0$ and $S_L = 1$, the “brake” on the right motor M_R is released, which makes the agent turn to the left. If light is perceived with both sensors, the agent releases both “brakes” and drives almost straight, slightly drifting to the right. In combination (Fig. 2 (B)), upon a presentation of light, these four behavioural modes lead to the following sequence of actions: 1.) A scanning turn to the right, until $S_L = 1$. 2.) A quick approach of the light from the right side. 3.) counter clockwise rotation around the light source. While the agent approaches the light source, it keeps bringing the light source in and out the sensory range of S_R (compare the rhythmically occurring drops of sensory and motor activity in Fig. 2 (C)). This strategy results in the chaining of nearly straight path segments in the approach trajectory, separated by punctual left turns (arrows in Fig. 2 (B)).

We now return to the agent’s value system. The estimator neuron M_5 outputs $E \approx 0$ if $S_L = 0$. The reason for this is that during the entire approach behaviour $S_L = 1$, and therefore $S_L = 0$ implies that the light has not yet been located, which only happens in the beginning of the trials if the agent is far away from the light source. During the nearly straight path segments, $S_L = S_R = 1$, which leads to $E \approx 0.5$, i.e. an intermediate estimate for an intermediate approach stage. While the agent cycles around the light source, $S_R = 0$ and $S_L = 1$, and the value system produces its maximum estimate, expressing that the light source has been reached. Notice also that the straight path segments which correspond to $E \approx 0.5$ become shorter as the agent comes closer to the light. Therefore, even though the value system has just three modes of output, its evolution over time can express a more gradual change in distance, if averaged over a time window: The average output increases with decreasing distance to the light.

Another event worth discussing in the trial depicted in Fig. 2 (B) and (C) occurs after the last displacement of the light source ($t > 2800$): As the displacement happens to bring the light source in the left visual field of the agent, it immediately enters the oscillating approach mode and its estimate therefore poorly corresponds to the actual distance measure which drops to 0. This dissonance can be seen as inevitable error due to the limited possibilities of the agent. However, we prefer to see it as superiority of the evolved estimator over the distance measure as a measure of performance: The comparably high output expresses the agent’s justified optimism to be at the light source soon. Such discrepancies between meaningful judgment signals generated by the agent and *a priori* specified performance measures were one of the key difficulties in designing the experiments. Even with the highly refined and complex fitness measure F_E (4), sometimes, “good” solutions in terms of the experimenter’s perception were replaced with less sophisticated ones by automated selection.

4.2 Value Guided Learning

Value systems are the proposed neural structures to guide ontogenetic adaptation. Can such mechanisms work if the value system is properly embodied? To investigate this question, we conducted another simple simulation experiment, in which the evolution of the robot controller is seen as the analogue of ontogenetic neural Darwinism as proposed in TNGS. The only parameters that evolve in this experiment are the strengths of the three synaptic connections from sensors to motors in the agent presented in the previous section (compare Fig. 2 (A)). The fitness measure F is substituted for the performance estimate $E(t)$. It is important to notice that in this set-up, the value system does not evolve, it just guides the evolutionary change of the synaptic weights to reinforce whatever behaviour leads to a high performance estimate $E(t)$.

Figure 3 (B) illustrates how with an embodied value system, value guided learning quickly results in a deterioration of light seeking behaviour, even though synaptic weights are just minimally altered. What the “value system” rewards is simply activation of the left light sensor but not the right. That this judgment

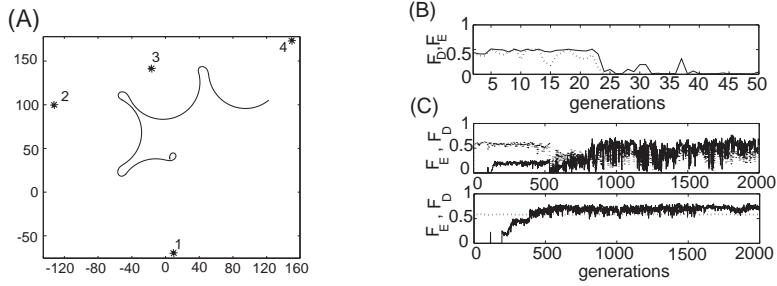


Fig. 3. (A) Light-avoiding trajectory of an agent after 50 generations of value guided learning. (B) The degeneration of light seeking performance F_D (solid line) and estimation performance F_E (dotted line) over time (same experiment) (C) Examples: F_E (solid) and F_D (dotted) in coevolving phototaxis (top) and fixed phototaxis (bottom)

means good light seeking behaviour during embodied interaction is a contribution of the sensorimotor context, and this meaning is removed if the system is functionally separated from the sensorimotor context. The gradual change of behaviour results in what we call “semantic drift” of the value system, i.e. the behaviour it rates as successful quickly ceases to be phototaxis (Fig. 3 (A)).

We see that the functional integration of the value system into the sensorimotor loop has far reaching consequences for the role this value system can play in the adaptation of behaviour dynamics. The reciprocal causal connections between behaviour generating system and value system undermine the idea of the value system as a top-down modulator of behaviour. But if the function of a neural structure whose activity we, as observers, can interpret as performance signal is not actually a value judgment, what could it be? This question is an open issue. One answer has already been given in Sect. 4.1 of this paper: Such a correspondence could be purely epiphenomenal and not bear any functional role in the generation of behaviour.

In an initial attempt to further investigate this question, we evolved agents in which the estimator neuron has the status of an interneuron and can project to other neurons. The most common structure we find in these networks is an excitatory self-connection in the estimator neuron that improves the estimation performance, but not phototaxis. In some of the networks that realise the same strategy described in Sect. 4.1, light seeking crucially depends on the activity of the estimator neuron. It serves to inhibit the right motor, as its activity is roughly in inverse correlation with the activity of the right sensor, and thereby takes part in inducing left turns if the light goes out of the right visual field. Its function is simply to relay and invert the right sensory signal. There is no end to the possible functions a “value system” could serve in the control of an embodied and situated agent. What the presented findings show is that the correspondence of neural activity to a behaviourally meaningful variable may well be plainly accidental.

4.3 The Evolution of Value Systems

Comparing the agents evolved to estimate value and seek lights to agents evolved to achieve just phototaxis (i.e. $F(i) = F_D(i)$), it turns out that the light seeking behaviour in agents that are evolved to estimate their performance is clearly suboptimal. Our first hypothesis to explain this phenomenon was a trade-off between the ability to perform judgments and the ability to find light quickly.

To test this hypothesis, we seeded evolution with successful light seeking agents and evolved combined light seeking and judgment behaviour on top, comparing conditions in which the sensorimotor behaviour was either fixed or continued to evolve with the value system. We expected the latter to be fitter, because the light seeking behaviour could be changed by evolutionary search to allow better estimation of performance. To our surprise, we found that both F_D and F_E were on average higher in the agents with fixed sensorimotor behaviour⁵. If good light seeking and good value estimation are possible at a time, why does the evolutionary search not find this solution? If we have a closer look at how the F_E and F_D component evolve in example evolutionary runs (Fig. 3, (C)), we see that the coevolutionary scenario (top) is much more noisy and good solutions repeatedly deteriorate. Apparently, in the presented set-up, a good estimation of the agent’s performance is very sensitive to behavioural noise and can only exceed a certain level if the sensorimotor coupling is completely fixed. This explains why value guided learning leads to such a rapid and devastating decay of behaviour: The noise sensitivity of value estimation accelerates semantic drift.

5 Discussion

Summarising the results from our simulation experiments, we presented an agent in which the capacity to judge on its level of performance with respect to a certain task crucially relies on the sensorimotor behaviour through which this task is realised. Without this sensorimotor context, the neural structure producing the performance estimate is meaningless, and if sensorimotor behaviour does not accommodate the need to estimate the level of performance, such judgment is only possible to a very limited degree, which renders the value system useless as internal supervisor of adaptive change.

Let us start our discussion by remembering the neural structures whose activity corresponds to salient events. From the presented results, two possible ways to interpret such structures result: a.) They could be embodied structures, integrated in a sensorimotor context, whose meaning has to be investigated and interpreted within this context and during situated interaction with an environment. b) They could be value systems that autonomously perform judgments about the significance of a situation and rewire the agent accordingly.

⁵ However, one of the seeded phototactic agents applies a strategy for phototaxis that does not seem to allow the estimation of performance. This suggests that there is at least some need for sensorimotor behaviour to accommodate judgment.

The presented results hopefully illustrate how these two options exclude each other: An “embodied value system” is a *contradictio in adjecto*. The existence of reciprocal causal links between value system and behaviour generating systems causes semantic drift of the value signal, which results in anarchy of development (see Sect. 4.2). But how could a value system not be embodied? Surely, we do not want to introduce magic meaning sensors or a magic master value system that ensures that the other value systems work smoothly. This smells too much of what Rutkowska calls “[b]uck passing to evolution” ([9], p. 292). If we struggle to explain the simple case without such scaffolding, the more abstract case will surely not become easier. The only way a value system architecture can work is a full embracement of the functional separation and pre-specification of meaning.

In the area of robotics, as shown in [12], we can design experiments rigidly enough to fixate meaning. But for an approach that aims at advancing past the stage of pre-specified motor programs, that refers to variable biomechanical properties in living organisms, the introduction of parts of the organism that are exempted from ontogeny, despite the constant material flux an organism undergoes, seems like a step backwards. It appears so inevitable that a random change would slightly change the context in which a value system is embedded, and the value-agnostic remainder of the organism would be unable to detect it or do anything about it. Furthermore, both in the area of biological modelling and in robotics, there is another unpleasant side-effect resulting from the introduction of disembodied and non-adaptive value systems: The impossibility of novel values. A rigid structure with *a priori* meaning can only work in situations that rely on phylogenetic constancies, the generation of new values in situations that our ancestors could not even have dreamt of asks for a different explanation.

We do not want to question that structures like the ones described as value systems exist in living organisms and that they play an important role in the adaptation of behaviour. In contrary, we think that the investigation of such mechanisms is important and intriguing. We plan follow-up experiments to the ones presented in Sect. 4.2, to investigate possible embodied functions that “neural value structures” could have for the adaptation of behaviour⁶. However, what we do want to question is that such components are or could be the loci of meaning. We question the idea that the generation of meaning can be separated functionally. Such components form part of an integrated system and their functionality both constrains and is constrained by this system they form part of, and therefore, they have to be interpreted as parts of a complex mechanism, not as encapsulated generators of judgment.

6 Conclusion

This paper does not have to be seen exclusively as a criticism of the value system as a locus of judgment, but as a general conceptual argument about correlation of neural activity with functional aspects of behaviour and how it does not

⁶ A crucial aspect to change is a task that requires long term adaptive modulation of behaviour, which was neither the case nor necessary in this paper.

entail, or even justify, the reduction of the respective function to the respective brain structure. Even though this point is not exactly novel, the enthusiasm with which researchers sympathetic to the embodied approach implement and develop “value system architectures”, in which a disembodied module is introduced to provide *a priori* specified criteria to guide embodied and situated lifetime development, provoked us to conduct the presented series of simple simulation experiments. These experiments illustrate the impossibility to reconcile functional reduction and the embodiment and situatedness of behaviour, which has been discussed in detail for the case of value system architectures, but extends to all models that feature a functional and structural separation of mechanisms of meaning generation from mechanisms of behaviour generation, i.e. all hybrid symbolic/embodied approaches to adaptive and intelligent behaviour: If a full-blown ghost in the machine has difficulties dealing with the variability of the external world, why would a vestigial ghost in the machine not face the same difficulties dealing with the variability of its bodily environment?

References

1. Beer, R. D.: *Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior*. In: P. Maes, M. J. Mataric, J.-A. Meyer, J. B. Pollack & S. W. Wilson (eds.): *From Animals to Animats 4. Proc. 4th Int. Conf. on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press 1996. 421–429.
2. Brooks, R. A., *Intelligence Without Reason*. Proceedings of 12th Int. Joint Conf. on Artificial Intelligence, Sydney, Australia, August 1991. pp. 569–595.
3. Edelman, G.: *The Remembered Present. A Biological Theory of Consciousness*. Basic Books, New York 1989.
4. Edelman, G.: *Neural Darwinism. The Theory of Neuronal Group Selection*. Oxford University Press, 1989.
5. Edelman, G.: *Naturalizing consciousness: a theoretical framework*. Proc Natl Acad Sci USA. Apr 29;100(9) 2003. pp 5520-5524.
6. Kandel, Eric R., James H. Schwartz and Thomas M. Jessel (eds.): *Principles of Neural Science*. Fourth Edition. McGraw & Hill, New York 2000.
7. Nolfi, S and D. Floreano: *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge MA 2000.
8. Pfeifer, R. and C. Scheier: *Understanding Intelligence*. MIT Press, Cambridge MA 1999.
9. Rutkowska, J.: *What’s value worth? Constraining Unsupervised Behaviour Acquisition*. In: Proc. of the Fourth European Conference on Artificial Life 1997. 290–298.
10. Sporns, O., and G.M. Edelman: *Solving Bernstein’s Problem: A Proposal for the Development of Coordinated Movement by Selection*. Child Dev. 64 (1993) 960–981.
11. Tuci, E., Quinn, M. and Harvey, I.: *Evolving fixed-weight networks for learning robots*. In Congress on Evolutionary Computation: CEC2002, IEEE Press 2002. pp 1970-1975.
12. Verschure, P., J. Wray, O. Sporns, G. Tononi and G.M. Edelman: *Multilevel analysis of classical conditioning in a behaving real world artifact*. Robotics and Autonomous Systems, 16 1995. 247-265.
13. Yamauchi, B. and Beer, R.D.: *Sequential behavior and learning in evolved dynamical neural networks*. Adaptive Behavior 2 (3) 1994:219-246.