From the Inside Looking Out: Self Extinguishing Perceptual Cues and the Constructed Worlds of Animats

Ian Macinnes and Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics,
John Maynard Smith Building, University of Sussex,
Falmer, Brighton, BN1 9QG
I.A.Macinnes@sussex.ac.uk, ezequiel@sussex.ac.uk

Abstract. Jakob von Uexküll's theory of the *Umwelt* is described and it is used to show how perceptual states can be defined. It is described how perceptual cues are selected over evolutionary time and defined by the organism that experiences them. It is then argued that by applying the model of the *Umwelt* to describe an animat's behaviour, we can model the normally distributed, dynamic activations of the animat as discrete perceptual states.

1 Introduction

This paper suggests we use a model to describe the behaviour of animats in terms of their constructed worlds and that by doing so, we stand to gain in two ways. The first way is that we find that have a model of the perceptual states of an animat, and so can apply various tools associated with state machines, including calculating the entropy of a constructed world. The second way is that by evolving the relationships between an animat and the objects in its environment that it interacts with, we are making the animat's constructed world more evolvable.

1.1 Controllers, Perceptions and the Constructed World

The nature of awareness and the mechanism of how perceptions are generated from physical matter has provoked much debate among philosophers. The most influential arguments regarding these subjects in the western world were made by the philosopher Immanuel Kant [1]. His arguments are complicated but among his conclusions is that the real nature of the world is unknowable, and we can only classify any possible external reality of an object through our perceptions of it rather than through direct knowledge of the object itself. Kant believed that the subjective sense of self is constructed through the process of the organisation

M. Capcarrere et al. (Eds.): ECAL 2005, LNAI 3630, pp. 11–20, 2005.

of perceptions. Whether perceptions are epiphenomena or play a causal role in behaviour, it is commonly supposed that a constructed world's function as part of an organism is managing information about the organism's environment and physical self.

The idea of using parameters to represent perceptions, mental states, or bits of knowledge has been highly influential for the majority of work performed in artificial intelligence and the disadvantages of it are well known [2]. It is tempting to use explicit representational states because of their clear information content. It is an easy task to calculate attributes such as the entropy of a controller with states over a period of time, or see how the states produce behaviour that result in other states [3]. Information and its transmission has long been used to explain animal behaviour, cognition, and evolution. It is very difficult to use these state based methods to calculate the information within the animats we evolve in evolutionary robotics as we do not use the concept of states, but instead allow control to self-organise from underlying parts. How are we able to determine that a state exists, or succeeds another state if they are not explicit? How do we assign states to the varying activation of a distributed structure?

Presently we analyse such systems in terms of dynamic processes such as attractors, limit cycles, etc. [4], together with the analysis of the role of specific neurons. The analysis required to understand even simple systems can be complicated and it is not clear how general the conclusions made about one system can be applied to another. There has been work performed in the field of artificial life on classifying information within an animat's distributed dynamic controller. Research has been performed on attempting to measure and analyse representational activations within an embodied robots neural network controller in terms of information [5]. Research has also been performed on the application of information theory to measure the information flow through perception-action loops of robots [6]. However, much work in evolutionary robotics tries to explain behaviour and perceptual organisation by analysing the dynamic activations of structural parts of a controller.

The line of enquiry described in this paper takes a different approach. It describes a model of the constructed world as defined by the ethologist Jakob von Uexküll [7], and why we gain by applying it to describe the behaviour of the animats we evolve. His theory was heavily influenced by the constructivist ideas of Kant. If we use his theory to explain animat's behaviour, then we have a simplified model of the constructed world generated by the animat. It is hoped that applying von Uexküll's theory will help overcome the problem of defining information within complex self-organising controllers, allow us to improve the evolvability of our animats, and perhaps lead to insights regarding the constructed worlds of robots and organisms.

First we will look at von Uexküll's theory of the *Umwelt*. Then we will see how we can distinguish perceptual states in his model. Finally we will argue that we gain by applying the model of the *Umwelt* to the less representative methodologies used in the field of artificial life.

2 Negotiating Functional Circles and the *Umwelt*

2.1 Functional Circles and an Organism's Constructed World

Von Uexküll explained purposeful animal behaviour by linking the organism's *phenomenal world* (the world as perceived) and its *effector world* (the world as enacted) into a single closed whole, the *Umwelt*. The *Umwelt* as defined by von Uexküll consists of a set of functional circles (figure 1). A functional circle defines the interaction between an organism and an element or object within its environment. They are abstract structures that tie together a subjective experience or perception (termed a *perceptual cue*) and the effect that the perception has on the behaviour of the organism (called an *effector cue*). Von Uexküll used the theory to demonstrate that an organism doesn't respond to its environment, but rather to its perception of the environment. Functional circles provide a model of how an organism's perceptual world is continually constructed as part of the organism's ongoing interaction with its environment. Using von Uexküll's model of the *Umwelt*, all of an organism's knowledge of itself and its environment is ultimately constructed from the perceptions within the subjective world it generates.

Von Uexküll himself provided the example of the purposeful behaviour of a female tick [7], of which he described three functional circles. The tick waits on a twig until a mammal moves close to the tick. The tick then jumps onto the mammal and burrows around in the mammal's fur until the tick finds a suitable place on its skin to bite a hole to suck blood from. This sequence of actions can be described using functional circles.

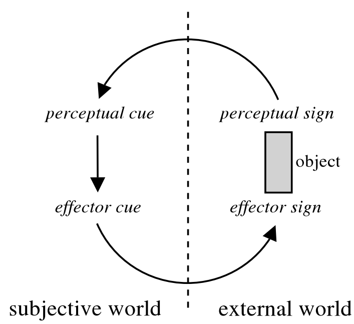

Fig. 1. A functional circle describes the abstract functional relationship between an organism and an object in its world. The perceptual sign of an object (its colour, shape, smell, or some more complex set of attributes) give rise to a perceptual cue which is defined as the subjective experience of the object in the organism's *Umwelt*. This leads to the creation of an effector cue which drives the animal to perform some action, thereby changing the organism's relationship to the object. A functional circle is an abstract description of a relationship, and it is hard to claim that a perceptual cue exists in any particular location, e.g. within a specific part of the nervous system.

First of all, the tick positions itself upon a suitable twig. If a mammal passes close by, the butyric acid emitted by the mammal provides a stimulus for the tick's smell receptor organ. This initiates a functional circle whose perceptual sign is the smell of the acid. A corresponding perceptual cue is produced, meaning that the mammal exists as an object in the tick's *Umwelt.* The resulting *effector cue* causes the tick to drop from the twig. The shock of the tick landing on the body of the mammal extinguishes the activation of the first functional circle. This means that the smell of the butyric acid no longer serves as a perceptual sign.

However, the shock of landing serves as the perceptual sign of a second functional circle. This initiates the behaviour in the tick of burrowing around in the mammal's fur until it encounters a patch of bare skin. The heat of the skin extinguishes the second functional circle, so the perceptual sign of the recent shock of the tick landing on the mammal no longer causes a perceptual cue and so the tick stops burrowing.

The third functional circle has the heat of the patch of skin serving as the perceptual sign. The heat produces a perceptual cue and so the patch of skin is experienced in the tick's *Umwelt.* The effector cue is produced and results in biting motions of the tick's mouth parts.

These actions can all be described as reflex behaviours with chemical or physical stimuli initiating fixed actions. Each behaviour has been selected by a process of natural selection to follow in the given order. We have used the model of von Uexküll's functional circles to explain them, which allow us to say that there is a perceptual cue associated with each circle.

If the organism is in a state suitable to allow the perceptual sign to serve as a stimulus, the perceptual sign causes the corresponding perceptual cue to be present within the constructed world of the organism. This is not to say that the perceptual sign *represents* some external object, or it is *presented* to the agent. As we shall see, the organism has chosen its own stimuli over its evolutionary history, and chooses whether or not to respond to it depending upon its current state. Both the stimulus, and the significance of the stimulus to the organism, is chosen by the organism. As a result, it may not be clear *what* the perceptual sign means to the organism.

A purely reactive organism relies on certain behaviours producing a predictable sequence of events in the environment that can be used as signs, e.g. an animal that hides in the dark depends upon the fact that the light intensity will decrease as it moves from the light thus extinguishing the stimulus. If the environment provided *all* the stimuli to cue the organism to produce its adapted behaviour, the organism would require only a simple and unsophisticated constructed world. However, a complex organism must engage in different sequences of actions at different times in an environment that provides the same potential stimuli. The organism must therefore choose which signs to respond to. A more complex organism is able to shape and vary how it perceives and hence responds to the world.

An organism with an *Umwelt* consisting of a small number of functional circles will perform only simple sequences of behaviour. This does not mean that the organism's physical environment is simple, but the environment *as perceived by the organism* is simple. The organism has selected over its evolutionary history the stimuli to which it should respond in order to engage with the world as reliably and efficiently as possible. Evolution acts to reduce the entropy of a constructed world despite of presence of environmental noise.

2.2 Self-extinguishing Functional Circles

Von Uexküll defined functional circles as always being self-extinguishing; that is, they produce behaviour in the organism that acts to remove the perceptual cue from the constructed world of the organism (figure 2). This can happen in two ways. The first way is that the functional circle causes the organism to alter its physical state to directly remove the stimulus, e.g. a photophobic organism will move away from a light, and thereby remove the stimulus that compels it to move. The second way is that one functional circle may *inhibit* the activation of another, e.g. the shock of the tick landing upon the mammal inhibits the first functional circle, even though the sign of the smell of butyric acid is still present.

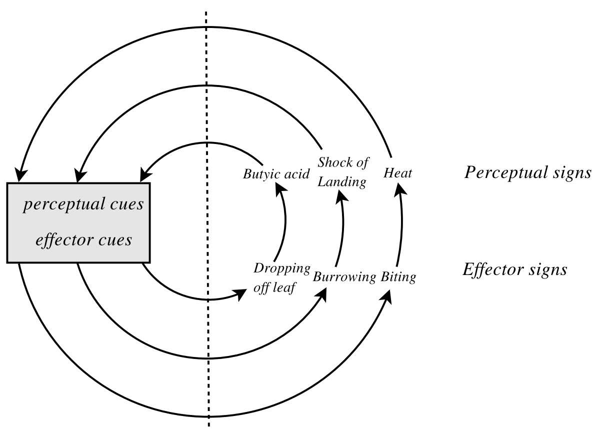

Fig. 2. The shaded region denotes that the functional circles negotiate among themselves as to which should be active. In the example of the behaviour of the tick, three functional circles negotiate. The perceptual cue of the first functional circle (i.e. the smell of the butyric acid), is inhibited and extinguished by the perceptual cue of the second functional circle (i.e. the shock of the tick landing upon the body of a mammal), and therefore no longer exists in the agents constructed world. This means that the butyric acid no longer acts as a stimulus, even though it may still be present, and the tick no longer perceives the smell of the acid. We can say that the two functional circles have negotiated to decide which should be active.

We can see that perceptual cues result in the creation of further perceptual cues, which act to extinguish the originator. The tick may either land on warm fur, or miss the mammal and hit the hard ground. The originating perceptual cue, the smell of the acid, invokes an action that results in one of two succeeding perceptual cues. Each extinguish the activation of the originator, and hence its associated behaviour or action.

2.3 An Organism Defines Its Own Perceptual Cues

A simple organism has evolved to respond to certain internal and external perceptual signs. The organism doesn't respond to a particular sign, e.g. the perception of heat, *because* it will get burnt, but because it has inherited a tendency to do so. The species to which an organism belongs has selected over evolutionary time the stimuli to which the organism should respond. The *meaning* of a perceptual sign may not be obvious to an outside observer, but it provides information that's meaningful to the organism that the organism uses to regulate or select its behaviour.

3 States of the Constructed World

3.1 It is Difficult to Classify States in Distributed Dynamic Controllers

It is necessary for an animat to co-ordinate its actions properly if its overall behaviour is to be purposeful. A simple behaviour is constructed from a number of discrete actions that occur one after another. The animat must know when to initiate which action, and at which point within the sequence, if the overall behaviour is to be successful. Traditional artificial intelligence techniques include the *finite state machine*. Using this technique, a distinct state and an action are associated together such that an animat in state A_s would always produce action A_a. Each state is linked to a limited number of other states. This reduces the possible sequences from action to action (i.e. state to state). The controller of such an animat is highly structured, and maps specific parts of the controller to specific local behaviours. This means modules that are responsible for producing individual behaviours can be programmed to react to specific stimuli and tested individually.

A continuous time recurrent neural network (CTRNN) controller is generally less formally structured and as a connectionist model has its activation states and any possible representations distributed throughout its structure. The state of each of its neurons can be altered by any other neuron it is connected to. This means it is not easy to claim a particular neuron or set of neurons of a CTRNN is responsible for a particular behaviour. CTRNN's use the idea of attractors, such that a CTRNN experiencing a particular attractor can be said to be in a particular state [8]. Desired actions are associated with different attractors. Attractors emerge from the dynamic activations of the CTRNN and are not explicitly defined. It is therefore difficult to isolate a behaviour or for a single neuron to selectively generate a specific behaviour. It is left to the controller as a whole to co-ordinate different behaviours within itself. This makes the robot's behaviour difficult to analyse and understand. It also makes it difficult to reuse existing behaviours in different contexts, because the production of a behaviour may depend upon the state of a number of neurons, which in turn depend upon the state of further neurons, and so on.

Although it is an ongoing and current research issue, there is currently no clear general method to classify the instantaneous activations of the components of a CTRNN into states. It is a complex task to analyse even simple artificial neural systems [8]. The alternative proposed in this paper is to explain behaviour in terms of functional circles as emergent structures since functional circles are a simple and direct model of perceptions and behaviour.

3.2 Perceptual Cues can be Modelled as States

A perceptual cue associated with a behaviour in the organism results in the activation of one or more succeeding perceptual cues with varying probabilities. The succeeding cues extinguish their originator using one of the two methods described above (see 2.2). Using this model, the constructed world acts as a manager and decides which perceptual signs should invoke actions and which should be ignored.

We can classify a state of the constructed world as the instantaneous activations of its perceptual cues. The state changes as the perceptual cues run in sequence as part of the process of the embodied animat engaging with the world. This state doesn't represent information about the environment. It is information about the *world as perceived* by the animat. When we apply this model, we are able to treat the perceptual world as a series of states that lead to other states, regulated by, and influencing the behaviour of, the embodied animat acting in its environment. The meaning of the perceptual cues, and therefore the information that defines the state of the constructed world, has been chosen and defined by the organism over evolutionary time.

3.3 Implementing an *Umwelt* for an Animat

The CTRNN was originally advocated as part of the methodology of minimal cognition [10]. Minimal cognition is an attempt to study the simplest form of cognition using the assumption that to understand a complex phenomena, we start with something we can understand and gradually increase its complexity. However, there is a distinct difference between the first, simplest animals, and the animats we evolve for our evolutionary robotics experiments. Minimal cognition was inspired by those simple animals with co-ordinated nervous systems. However, it is likely that the behaviour of todays simple animals originated early in evolutionary time from simple, reactive, independent but co-operating reflexes (figure 4). The behaviour of simple animals is co-ordinated as they are comprised of sequences of actions that evolved to occur in a meaningful order. This doesn't

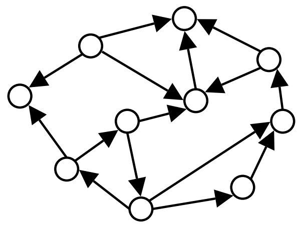

Fig. 3. It is difficult adding to an evolved robot's repertoire of behaviour because it requires integrating additional sensorimotor loops to an existing connectionist controller that defines preexisting behaviours. Sensorimotor loops in robots are partly realised as dynamic states along components of the controller. Some of these parts may be common to the action of several sensorimotor loops. A mutation to one of these common parts may be beneficial to one behaviour but detrimental to another; this means that sensorimotor loops from which each desired behaviour is composed can not evolve in isolation from each other, and this may constrain the evolution of both. Subsumption architecture [9] was developed in behaviour based robotics partly to solve the problem of scalability but no generally accepted methodology exists for evolutionary robotics.

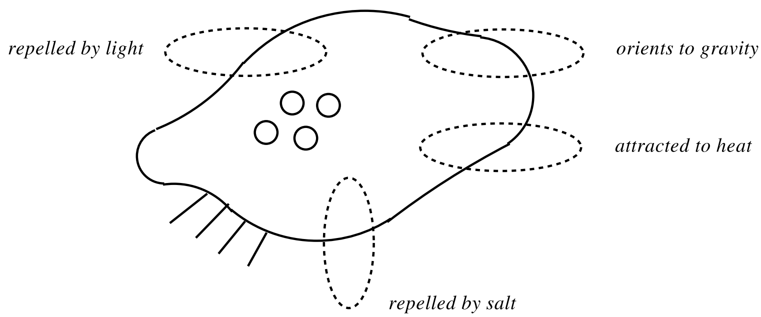

Fig. 4. Simple animals do not have a controller but operate through reflex reactions that are co-ordinated by the environment. Over evolutionary time, the organism may evolve more complex behaviour by negotiating and therefore selecting which reflexes to respond to.

mean that the environment is the controller of the organism. The agent within its environment reacts in a purposeful way to individual stimuli that it expects to occur in a certain order. The environment has to provide these stimuli in a given order for the agent to function properly.

Organisms without explicit controllers make up almost all living things, and yet they still manage to display adaptive behaviour. The functional circle model provides a union between the sequencing of actions and the actions themselves. Together with evolving perceptual cues, evolving negotiating functional circles allows agents to chose their own stimuli, and alter and reuse sequences of actions without altering the actions themselves. We can at some later point derive a co-ordinated control system developed from low-level reflexes.

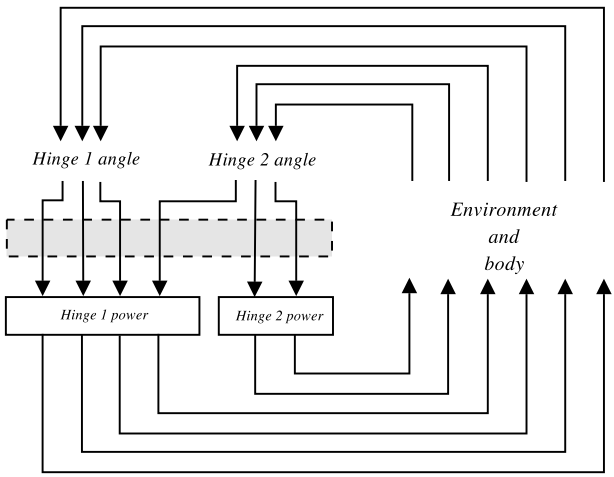

Fig. 5. Although functional circles can negotiate through the environment, evolved functional circles involving input sensors are likely to negotiate at the most mutable point, inside the controller. We should see the controller as mostly a negotiator of functional circles. We can view them as negotiating between perceptual cues.

It may not be possible to directly implement a model of functional circles (indeed every responding system can be described in terms of functional circles) but only to use mechanisms that encourage and highlight their essential properties. We can explain the behaviour of the emergent structures of CTRNN's in terms of functional circles. We often look at activated neurons that correspond with certain behaviours [8]. But we can alter our models to encourage the production of perceptual cues (figure 5).

The perceptual signs were clear and easily identifiable in the case of the tick. However, we do not need to know what perceptual signs the agent is reacting to if we chose a model whereby we can identify, by the nature of the model, its perceptual cues. We do not need to know or understand what the perceptual cue is actually reacting to. If we treat the controller as a negotiating part, and allow our perceptual cues to evolve separately, then we will be able to use this model.

The functional circle model can be applied to any the behaviour of any animat, including animats controlled by CTRNN's, even though the representation of their perceptual cues may be distributed across a network. All controllers can be described in functional circle terms because the *Umwelt* is an abstract model of behaviour and subjective experience, and how these affect the outside world.

4 Conclusion

Von Uexküll's original model of the *Umwelt* provides a model of an organism's constructed world. The constructed world is the environment as perceived by the organism. This is an abstract model that can be used to describe the behaviour of any animat, regardless of how their controlling mechanisms are implemented. If we apply this model to describe the behaviour of the CTRNN controllers of our evolved animats, we benefit in two ways. The first is that by treating functional circles as emerged structures of CTRNN controllers and evolving the manner in which they negotiate together, we are evolving the relationships of an animat with its environment. The second is that we can define varying states for the constructed world of the animat. We can apply the tools associated with state machines while enjoying the benefits of a non-representational controller.

References

1. Kant, I.: A Critique of Pure Reason. 1930 edn. G. Bell and Sons, Ltd., London (1781) English Translation.
2. Brooks, R.: Intelligence without reason. In Myopoulos, J., Reiter, R., eds.: Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91), Sydney, Australia, Morgan Kaufmann publishers Inc.: San Mateo, CA (1991) 569–595
3. Ashby, W.: An Introduction to Cybernetics. John Wiley & Sons, Inc., New York (1956)
4. Strogatz, S.: Nonlinear Dynamics and Chaos. 2000 edn. Westview Press, Cambridge,MA (1994)
5. Avraham, H., Chechik, G., Ruppin, E.: Are there representations in embodied evolved agents? taking measures. In W.Banzhaf, T.Christaller, P.Dittrich, J.T.Kim, J.Ziegler, eds.: Advances in Artificial Life - Proceedings of the 7th European Conference on Artificial Life (ECAL). Volume 2801 of Lecture Notes in Artificial Intelligence., Springer Verlag Berlin, Heidelberg (2003) 743–752
6. Klyubin, A., Polani, D., Nehaniv, C.: Tracking information flow through the environment: Simple cases of stigmergy. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., Watson, R., eds.: Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, Cambridge, Massachusetts, MIT Press (2004) 562–568
7. von Uexküll, J.: A stroll through the worlds of animals and men : A picture book of invisible worlds. In Schiller, C., ed.: Instinctive Behaviour: The development of a modern concept. 1957 edn. International Universities Press, Inc., New York (1934) English Translation.
8. Beer, R.: The dynamics of active categorical perception in an evolved model agent. Adaptive Behaviour **4** (2003) 209–243
9. Brooks, R.: Intelligence without representation. In Meyer, A.R., Guttag, J.V., Rivest, R.L., Szolovits, P., eds.: Research Directions in Computer Science: An MIT Perspective. MIT Press, Cambridge, MA (1991) 249–276
10. Beer, R.: The dynamics of adaptive behaviour: A research program. Robotics and Autonomous Systems **20** (1997) 765–776