How Cognizers Come to Know their World and What this Implies for Cognitive Robots

Alexander I. Kovács **Haruki Ueno**

The Graduate University for Advanced Studies
National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-ku, 101-8430 Tokyo, Japan
{aik@grad. | ueno@}nii.ac.jp

Abstract

This paper points out some fundamental issues of particular relevance to the study of artificial cognitive development and the fabrication of cognizers (that is cognitive systems), notably: (a) the way a cognizer has access to its environment, (b) the fundamental difference between our (the scientist's) models and those internal to a cognizer, (c) and the contingency of cognition. These are not separate but intertwined issues and will be discussed as such. We start by introducing the notion of a *meter* and use it to generalize Uexküll's *umwelt* concept. Clarifying how any cognizer, in principle, can get to know about its environment we arrive at the notion of a cognizer's *umwelt*. This leads to the question of what and how a cognizer can learn about its environment. Since this inevitably involves the notions *model* and *representation* we discuss how we must carefully distinguish between our own formal (encodingist) models and the interactive ones internal to a cognizer, which necessitate the use of oscillators. In addition to environment as unspecified source observables (vs. cognizer) and *umwelt*, we also work out the concepts of a cognizer's *cognitive body* and *cognitive substrate*. In this framework we formulate the principle of the Contingency of Cognition, which states that cognition depends on the combination of all four concepts. We finally work out a number of ramifications pertinent to the fabrication of (possibly veritable) cognizers such as cognitive robots.

Introduction

Once we think about endowing robots with "real cognition" (whatever that means) a number of fundamental and programmatic issues are raised that need to be addressed before we think about a robot's tasks, knowledge or behavior. For one thing, the question is open (1) whether *veritable* artificial cognizers are possible; or the question (2) whether cognition presupposes aliveness (whatever that means)—we leave these to philosophers to decide, because the very nature of the fields like developmental or cognitive robotics requires that we be faithful at this point in time, i.e. that the answers be (1) yes and (2) no.

If by cognitive development we mean that a system comes to know new things about its environment, integrates those with what it already knows, and utilizes that knowledge, the question is: How should we go about artificially reproducing this in a robot? The current state of the fields concerned with cognitive functions in real or artificial cognizers is however such that at present no straight-forward answer is possible. All these fields ought to be based on a *general* Science of Cognitive Systems; a foundational science that doesn't yet exist. Such a science is necessary, for we need to understand what it is, *in general*, about systems, which makes them such that they can be said to "possess" what Descartes called the *cogito*. *Without such an understanding we cannot see how we are to successfully fabricate cognitively developing artifacts, such as the machines we currently call robots*.

What we perceive to be one of the fundamental questions that need to be understood, is *how* any cognizer, artificial or not, comes to know about its environment. We shall start with this question. After that we need to address the issue of models and how they enter into the study of cognizers. In later sections we will then introduce the concepts of *cognitive body* and *substrate* and formulate a principle, which is true for any cognizer, namely the contingency of its cognitive capabilities. We will then, in the final part of the paper, point towards a number of ramifications regarding cognitive robots which follow from our framework and the contingency of cognition.

Meters

To drive at the notion of a meter, let us suppose we have two natural systems, S_1 and S_2. Informally, we identify S_1 and S_2 as natural systems because both possess certain qualities that we can perceive and certain relations (not directly perceivable) among those percepts that we attribute to the systems. Or in the words of Rosen (1985, p. 47), a "natural system is a set of qualities, to which definite relations can be imputed." Let us refer to perceptible qualities as *observables*; and let us call relations between observables *linkages*. These "observables are the fundamental units of natural systems, just as percepts are the fundamental units of [human] experience." (ibid.) The sources of percepts are qualities, or observables, of natural systems. If now the interaction between S_1 and S_2 causes some perceivable change, say in S_2, the cause must lie in an observable of S_1, and vice versa. Rosen (1985, p. 48) thus takes "observables *in general* as the vehicles through which interactions between natural systems occur, and which are responsible for the ultimately perceptible changes in the interacting systems arising from the interaction."

We can now choose an *arbitrary* natural system M, which we call a *meter*; Rosen (1978, 1985, 1991) employs the device of a *meter* to generalize what we could call sensory impressions, or percepts; we shall use it later to generalize the umwelt concept of Uexküll. Now, by placing M in interaction with arbitrary systems S, some of those interactions will result in a perceptible change in M, which may be different for different S. M thus *defines* an observable (namely those qualities of natural systems that can impose a change in M) and we can say that this observable takes on different values (the different perceivable resultant states of M). The difference in the changes arising from interaction with different S can serve as a measure for that value (hence the term "meter"). Note explicitly that we have nowhere assumed that these values need to be encodeable as numbers.

In order to formalize the above, Rosen (1978, 1985) introduces (a) S_A, the *set of abstract states* of a natural system S; and (b) the *spectrum*, the set of all possible values, which a meter M may assume. Roughly speaking, "an abstract state of a system is that on which the observables of the system assume their values. In this sense, the value of any observable on an abstract state is a way of naming that state; i.e. is already an encoding of the abstract state." (Rosen 1985, p. 126) Both concepts are most peculiar in that they do not lie completely in either the external nor formal, or platonic world. This is indicated in the following figure by the dashed line. S_A as a set is part of the formal world; its members however are part of the external world and the same is true for the *spectrum*.

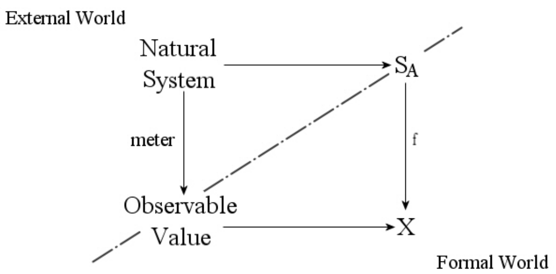

Figure 1 Formalization of measurement with meter (adapted from Rosen 1985, p. 127).

In summary, we associate with every natural system a set of abstract states S_A which we learn about through acts of measurement. By encoding the *meter* as a function f: $S_A \rightarrow X$, this is formalized as f assuming a name in X on an abstract state. The function f thus represents the observable defined by M. And since a *meter* M, by its *spectrum*, defines an observable, the choice of M determines what we learn about a natural system S.

As we will proceed to show, the *meter* is a most fruitful conceptual place to start our investigation; at the same time, however, it is the place where we can illustrate in its most generality the *encodingist dilemma* of which we will have to say more below. For now, we just point out that in above discussion we always talked of "perceptible change" in a *meter* and other natural systems. We retained an all−mighty observer, which the figure hides. As a matter of fact, *the observer is represented in the figure as the dashed line.* An external observer is necessary to associate, or encode meter values with labels in the set X, that is, to evaluate function f. In fact, the observer must make *another* measurement (e.g. by looking at the meter), to find out the value. The concept of a *meter* thus only makes sense if we have recourse to an observer. Rosen was not careless: he assumed this observer in the scientist. Before coming back to meters, the next section will introduce another circle of ideas, which also originates with Rosen.

Models and Abstraction

All disciplines involved in the study of cognition universally agree that a cognizer *must* posses (in some form or another) models, that is representations of its environment; cognitive development is thus essentially concerned with the development of such models, their storage and usage in the cognizer. In order to fix ideas we will directly introduce Rosen's *modeling relationship* (Rosen 1985) rather than start a discussion of the 20 something definitions in the Oxford English Dictionary, of what a model is.

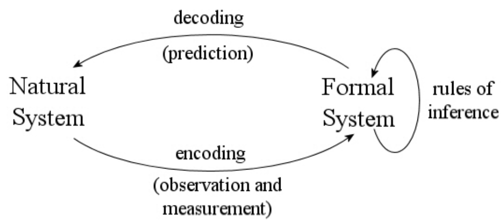

Figure 2 The modeling relationship (adapted from Rosen 1985, p. 74).

We can say that a modeling relationship obtains between a natural system and a formal system if the entailment structure in the formal system (inferences) mirrors the entailment structure of the natural system (causation); we then also say that the formal system is a *model* of the natural system. By measuring observables of natural systems we can come to know linkages obtaining between those observables; and by encoding the observables and their values these linkages can be expressed as formal systems. Rosen (1985) stresses, that both the encoding and decoding functions neither belong to the natural or formal system nor are they in any way entailed by them; the real creative act in modeling lies specifically in finding those two functions.

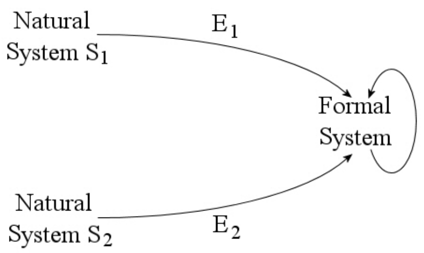

Figure 3 Analogy between two natural systems (adapted from Rosen 1985, p. 77).

A special situation is depicted in Figure 3. There, two different encodings E_1 and E_2 encode two different natural systems into the same formal system, such that the formal system is a model of both S_1 and S_2. In this case we say that S_1 and S_2 are *analogs*; each can be regarded as a model of the other. Since formal models are denizens of a platonic world, what is actually meant (but seldom explicated) by "a cognizer possessing a model" is captured by this characterization of analogy. However, analogy, in this Rosenean sense, is a relationship between *natural systems* and *relative to the encodings* E_1 and E_2; so, strictly speaking, to say S_1 and S_2 are models of each other is an abuse of language (Rosen 1985, p. 78). The functional composition $E_2^{-1}E_1$ (that is traversing in Figure 3 from S_1 via the formal system to S_2) represents an encoding of states of S_1 into states of S_2 (idid., p. 340). If a cognizer possesses a model in this sense of analogy, we shall call this model *encodingist*. One thing encodingist models assume is that such encodings can be found; we shall argue that this assumption is wrong and hence that models contained in cognizers must be fundamentally different from encodingist models.

Let's move on to *abstraction*. Countering the intuition abstraction belonged to the realm of theorizing only, Rosen (1985) shows that it already pertains to the very notion of a *meter* (and is thus relevant to experimental fields as well); this innocent insight is relevant to us, as we will see shortly. Given a meter M, let f be the observable associated with it. Of all the possible qualities of a system S, the only one M measure is f. The only thing about S visible to M, in principles, is its f-ness. That is, M *abstracts* away from S everything, but f. There is a second aspect in which M abstracts: Namely, if for two different abstract states s_1 and s_2 we have $f(s_1)=f(s_2)$ then M cannot distinguish between them. For the meter M, the natural system S is accessible only in this abstraction. This may not be how S looks to other meters at all.

But *this is precisely the content of Uexküll's concept of the umwelt of an animal*, which we have just generalized to the "Rosen" meter as promised. With reference to Kantian philosophy, Uexküll (1973/1928) ascribed to every subject (animal or human) its *umwelt* (from German *ambience* or *environment*), meaning the environment as it *appears* to the subject. But as we have just seen, to a single "Rosen" meter no other qualities about a system exist but those embodied by the observable it measures; it thus makes sense to speak of the *umwelt* of a meter.

From above discussion pertaining to abstraction it is now immediate that any formal system F can be a model only of a subsystem S' of a natural system S. F can say nothing about observables not encoded in it or any linkages that obtain among those non-encoded observables or linkages between encoded and non-encoded observables. If every observation is an act of abstraction, neglecting other qualities present in the external world, then the qualities captured by those abstractions can refer only to restricted parts of the world. Based on considerations like these Rosen can prove his *Main Theorem* (1985, p. 293), which *roughly* states the following: The trajectories of two systems S and S'—both identical except that S is opened up to an additional modality of environmental interaction to which S' is closed—will diverge if both systems are started in a common initial state. Further, Rosen extends this result to show that—as a result of abstraction—the same discrepancy arises between the predictions of a model and the dynamic behavior of the system being modeled. In other words, there will always be a critical instant for which the modeling diagram will no longer *approximately commute* (defined by Rosen) anymore. The simplest solution to make the diagram commute again, Rosen points out, is to *recalibrate the model*, "by repeating this process [of recalibration] often enough, even a faulty model can retain approximate commutativity indefinitely". (1985, p. 304).

Interactive Models and Oscillators

We mentioned above that the relation $E_2^{-1}E_1$ in Figure 3 represents an encoding of states of S_1 into states of S_2, where S_1 was said to be a model of S_2. Suppose now we have a natural system S that contains S_1 as a subsystem. It is immediate that to S its subsystem S_1 cannot represent S_2 by *correspondence via such encodings* for they are constructions of an outside observer; we have seen this in its most general form in our discussion of the formalization of the Rosen meter and in Figure 1. This situation is precisely the content of the critique of Bickhard and Terveen (1995) which can be stated succinctly thus: any attempt to explain mental representation by correspondence via encodings is deemed to circularity; it presupposes what it is trying to explain. A moment's thought reveals that this must be so for any cognizer at all, not only animals.

But if not via correspondence, how else are internal models to be about anything? Bickhard and Terveen offer an *interactive* conception of representation; we shall therefore call internal models of cognizers *interactive models*. Timing, so Bickhard and Terveen (1995, p. 84), is *foundational* to representation, and van Gelder and Port

(1995, p. 19) equivocate: in cognition "*timing* always matters". Timing requires oscillators. Though van Gelder and Port do stress the timing issue, oscillators are not part of the Dynamical Hypothesis [1]; the essential role of oscillators is most directly enunciated by Bickhard. We will employ *oscillators* as a conceptual device in the sequel. The term will gradually be filled with meaning below.

From our discussion of Rosen's modeling relationship, however, we see that *our* (the scientist's) formal models *must necessarily* be encodingist; this is how we do science: we observe, measure, and encode into formal systems. We now see that *our* formal models must be *fundamentally* different from those internal to any cognizer (and, a fortiori, those presumably in our own brains) because those cannot be based on encodings; the two must therefore be carefully distinguished. Further, we conjecture that it will be impossible to find the encodings necessary to show that an interactive model is analogous in Rosen's sense to another natural system. An interactive model just doesn't contain the necessary referents. This is so because there will be no discreet functional relationships between the entities embodying the models.

The Cognitive Body and Cognitive Substrate

We will now bring together the concepts of *meter*, *oscillator*, and interactive models in order to work out their fundamental role in cognition, in how any cognizer comes to know its environment.

Though there is nothing in the state of a meter that announces any correspondence to the observable that gave rise to it, it is correct to say that it somehow *embodies* the observable, for it defines it in the first place. We will also say that a meter *transduces* an observable into its own state. As far as a cognizer is concerned, there are a priori no natural systems S but just an unspecified *environment* E which it attempts to know about; there is no better way to put this: "S is initially *unknown*, veiled completely in its noumenal and phenomenal shrouds." (Rosen 1991, p. 157, original emphasis).

So, *in order to come to know its environment any cognizer must interact with it with meters*. As pointed out above, by fixing a certain meter, we fix the possible modes of interaction with the environment and thus an umwelt, a way the environment looks to the meter. By definition, the more *different* meters a cognizer uses to interact with the environment, the more observables it will have access to. Just having many different meters is however not enough. The cognizer must detect regularities between the observables it has access to, in order to discover linkages between observables. Without discovering any linkages, it cannot come to possess models.

[1] The Dynamical Hypothesis states: "Natural cognitive systems are dynamical systems, and are best understood from the perspective of dynamics." (van Gelder & Port, 1995, p. 5)

So, not only the number of meters and their respective *spectra*, but also their arrangement relative to each other is important. While some usual functions of bodies are perhaps to keep certain matter out and other matter in and maybe that of protection and providing structural stability, the *cognitive body* of a cognizer has (1) the function of making possible interactions with the environment such that external observables can change states in meters and (2) determine certain arrangements of those meters. If the reader dislikes the term meter, she may substitute it with receptor, or sensor.

To make these ideas more concrete, let's start with two meters, M_1 and M_2. We can arrange for M_1 to transduce an observable of the environment and M_2 an observable of M_1 (see Figure 4a below). Any meter, qua natural system, possesses itself qualities which can in turn be the source of change in other meters. The case interesting us most, of course, is when the qualities of M_1 are somehow related to the observable embodied in M_1 and thus ultimately to S_A.

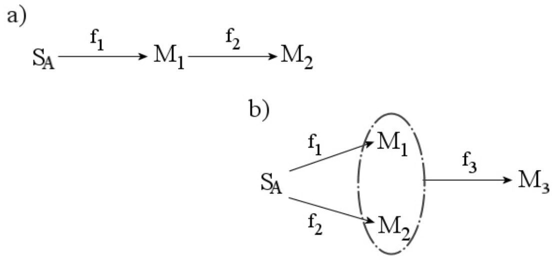

Figure 4 Different arrangements of meters.

The second case, depicted in Figure 4b, is also most interesting. Meters M_1 and M_2 independently interact with the environment and embody *possibly different* observables. If M_1 and M_2 are now such that both transduce their respective observables into the same observable quality, M_3, interacting with both M_1 and M_2 as one system, actually transduces a relationship between the observables f_1 and f_2. In this case we shall say that M_1 and M_2 transduce into a *common currency* (the observable defined by M_3). Even though relationships are not directly perceivable, networks of meters can come to embody them.

Continuing this line of thought we see that we can imagine arbitrary networks of meters. Such networks are comprised of two different kinds of meters. The first (such as M_1 in Figure 4) interacts directly with the environment and thus arbitrary observables. This kind is part of the concept of cognitive body. The second (such as M_3) are those that transduce the *common currency*.

Any oscillator, again qua natural system, is a meter. The term oscillator now stresses the fact that we are not interested in "implicit differentiators" (as the term *meter* somewhat connotes), but "indications of potentialities for interaction" (Bickhard & Terveen, 1995). Oscillators can do this by oscillating, their oscillation being modulated by the transduction process. If we now have a number of oscillators such that all can modulate each other ("use" a *common currency*) we shall call such a network of oscillators a *cognitive substrate*. Oscillators, like all meters, embody observables in their state, which are oscillations. Interactive models, we see now, are embodied in the action and interaction of oscillators in the cognitive substrate of a cognizer. The meaning of interactive model is thus that "representational content of an interactive representation is constituted in the organization of such webs or indications" (Bickhard and Terveen 1995, p. 305). Cognitive substrates can thus embody new features or relations that are not directly part of a cognizers umwelt but which it imputes back to its umwelt.

It was Uexküll's (1973) great insight to grant each subject its own specific *umwelt*; that is to note that there are as many umwelten[2] as there are subjects. We see now how the cognitive body together with the cognitive substrate of a cognizer determines possible modes of interaction with the environment as well as the linkages that can be discovered. It is in this way that the *combination* of the umwelten of all meters involved comes to bear upon a cognizer as a whole—to any cognizer, that is, not just animals, the subjects of Uexküll.

What we call *cognitive body* might appear to be expressed by the terms *morphology* or *sensor-motor system* (cf. Pfeifer & Scheier, 1995); and what we call *cognitive substrate* seems identical to what is variously called *controller* or *developmental program* (Weng and Zhang, 2002). However, this new terminology is an attempt at complete generality, i.e. supposed to pertain to *any cognizer* at all not just animals or cognitive robots. By way of example, consider the human body, which contains a nervous system and an immune system (also see the conclusion for this point). In our terminology it can now be stated in a precise way that the human body contains *two different cognitive substrates*, each with its own separate *cognitive body*. Nobody would however say that the immune system is a *controller*, nor is it obvious what the *morphology* of the immune system should be. In the same vein, we speak of *cognitive* development rather than, say, *mental development* (Weng and Zhang, 2002); "mental" carries strong connotations relating to the minds of organisms. The term cognitive seems to be more generally applicable.

The Contingency of Cognition

In this section we will try to formulate the principle of the contingency of cognition. This principle, too, is supposed to be general in application and thus to apply to cognitively developing robots as a special case.

What we have hoped to arrive at is that no matter what the particular nature of the meters in the cognitive body and the oscillators in the cognitive substrate, the framework allows us to talk about what is minimally needed to characterize a cognizer. So, when we try to determine what it is about a system that makes it cognitive, we ought to be able to point to subsystems that realize our notions of *cognitive body* and *cognitive substrate*. Specifically, it claims the *existence* of oscillators in the substrate of cognition—predicts their *necessity*. Conversely, how can we fabricate a cognizer? Well, cognitively embody it properly and give it a cognitive substrate. For the body determines the way a cognizer can come to know about its environment (fixes observables) and the substrate what linkages between them can be discovered: *together they determine the models the cognizer can come to possess*. The rest must come from interaction with the environment.

If only we can agree that at the root of cognition lies the ability of a system to come to posses models of its environment, then almost by definition we see that any change in the combination of *environment*, *umwelt*, *cognitive body* and *cognitive substrate* will change the possible interactive models that can be developed and thus all cognitive abilities of the cognizer. Metaphorically speaking, the cognizer gets answers only to the questions it asks and the answers will only be those the environment can give. To come to know the environment, the cognizer with its (cognitive) body and substrate asks the question that the environment must answer. *It is the specific combination of environment, umwelt, body and substrate that determine the resultant cognitive performance of any cognizer; not any of them in isolation.* Specifically, the principle says that cognition is not something *intrinsic* to the animal brain or any developmental program (Weng and Zhang, 2002). Descriptions of all interacting parts must figure *equally* in order to understand the capabilities of a cognizer.

Implications for Cognitive Robots

In this section we shall itemize a number of implications of the framework developed so far.

(1) We have argued that any cognizer comes to know its environment only if different observables of the environment can impose a change of state in appropriate meters and that the only way to *autonomously* discover linkages between observables is by special meters (oscillators) being connected in networks. So, the first requirement for any autonomous cognitive development is to *outfit artificial cognizers (such as cognitive robots) with a proper cognitive body and substrate*.

(2) It follows from the contingency of cognition that we must *adjust our expectations* of cognitive performance according to all parts entering in the equation. Specifically, an entity with an impoverished umwelt cannot be expected to be a cognitive high-flyer: *no amount of computation* can recover any linkages in the environment if they are hidden by the umwelt, that is, if no linkages exist with accessible observables in the umwelt of a cognizer. So, for a cognitive robot it is important to endow it with as rich an umwelt as possible.

(3) The scientist's formal models are fundamentally different from those internal to the cognizer, which are interactive in the Bickhardian sense and must be embodied

[2] German plural of *umwelt*.

by a cognitive substrate. That doesn't mean that we cannot computationally simulate the cognitive substrate, it just means that it will be meaningless to search for any referents in the substrate, which encode anything in the environment of the cognizer. In other words: setting up any kind of building block construction of representation will not work. The only material available to a cognizer (real or artificial) are observables and linkages embodied in the (inter)activity of oscillators in its cognitive substrate. *This is the stuff autonomous mental development* must be based on—in any cognizer, a fortiori in humans as well as in robots.

(4) Since any computer simulation is an abstraction (it only contains what we encode into it) it follows from Rosen's main theorem that the simulation will diverge from what it simulates. Together with our principle of the contingency of cognition this means that if we hope to evolve cognizers by some form of *simulated evolution* in which all parts on which cognition is contingent must be simulated we should be prepared to evolve cognizers that will not work when put into real environments. By "will not work" we mean that they will have evolved in *umwelten* so different from the real world, that their internal models will be too much out of sync. In other words, *cognitive development for real world cognizers must happen in the real world*. We can appreciate the severity of this problem, once we realize that we can simulate only rather simple environments.

Conclusion

It is up to a cognizer, artificial or not, to carve up the external world into systems. The *umwelt* concept, already an issue in as unspecified a system as a *meter*, percolates all the way up to networks of meters, however complex these networks may be; and thus reminds us of the fact that different cognizers will necessarily carve up the world differently for they will ultimately have access to different observables and have different capabilities of noticing linkages between them. It is from this observation that any discussion of cognitive development must start; this is so because cognitive development fundamentally means to acquire knowledge about the environment and the acquirable models are at the root of *all* cognitive performance of a cognizer. Discovering models presupposes two things: observables and linkages; it is the *cognitive body* that lets the cognizer "see" observables and it is the *cognitive substrate* with which linkages are "created" and imputed back to the environment. But it is the environment which must impose dynamic changes in the meters of the cognizer. The substrate not only embodies the models it is also directly involved with their creation. It is the difference in those models that leads to utterly different cognitive capabilities, and since the creation of models depends on cognitive substrate, cognitive body, umwelt, and environment, a fortiori all of cognition becomes contingent.

In order for the young field of developmental robotics (Weng and Zhang, 2002) to make progress we believe it should be based on a general and theoretical inquiry into cognition and cognitive systems. Such a science must necessarily take a much more comprehensive look at cognition than neuroscience, cognitive science, or artificial intelligence. In this paper we have avoided reference to any neurons or neural networks and the formal models thereof; the reason being that the animal nervous system is but one instance of a cognitive substrate known to us: The immune system is often referred to as being a cognitive system and in fact all systems known under the rubric of complex adaptive systems are tacitly ascribed cognitive capacities. What needs to be explained is cognition in all these manifestations.

Our framework hints at a possibly very fruitful formalization of the terms *embodiment* and *situatedness*; and further a formalization of Uexküll's *umwelt* concept, which could lead to a measure of "degree of embodiment" that could be put into relation to some index (such as the size) of the *cognitive substrate*. This quotient could perhaps be used to explain differences in cognitive capabilities between species of animals. We leave it for future work to explore these ideas.

Acknowledgements

AIK was supported by the Joint Studies Program of The Graduate University for Advanced Studies and an NII Fellowship Grant kindly provided by Sumitomo Electric Industries, Ltd. We thank Sébastien Duval and Eric Platon of NII and the anonymous reviewers for critical reading of the manuscript.

References

Bickhard, M., and Terveen, L. 1995. *Foundational Issues in Arti-ficial Intelligence and Cognitive Science*. New York: Elsevier Science.

Pfeifer, R., and Scheier, C. 1999. *Understanding Intelligence*. Cambridge, Mass.: MIT Press.

Rosen, R. 1978. *Fundamentals of Measurement and Represen-tation of Natural Systems*. New York: Elsevier Science.

Rosen, R. 1985. *Anticipatory Systems*. Oxford: Pergamon Press.

Rosen, R. 1991. *Life Itself.* New York: Columbia University Press.

Uexküll, J. von. 1973. *Theoretische Biologie [Theoretical Biology]*. Frankfurt a. M.: Suhrkamp. Originally published 1928.

Van Gelder, T., and Port, R. 1995. It's About Time. In Port, R., and van Gelder, T., eds. Mind as Motion, pp. 1-43. Cambridge, Mass.: MIT Press.

Weng, J., and Zhang, Y. 2002. Developmental Robots: A New Paradigm. In Prince, C., et al., eds. *Proc. 2nd Int. Workshop on Epigenetic Robotics*, pp. 163-174, Edinburgh, Scotland.