UMWELTS and Artificial Devices

A Reflection on the text of Claus Emeche :
Does a robot have an Umwelt ?

Ricardo R. Gudwin
DCA-FEEC-UNICAMP
gudwin@dca.fee.unicamp.br

Introduction

In the text "Does a robot have an Umwelt ? Reflections on the Qualitative Biosemiotics of Jakob von Uexkull" (Emeche, 1999), Claus Emeche frames a very important question that has been around the scientific community for a while : can artificial devices be somewhat compared to living organisms ?

Despite Dr. Emeche frames the problem from the point of view of Umwelts and robots, this is the basic question that is behind his argumentation. There seems to be a sense of discomfort when we try to apply concepts derived from biology in order to include artificially constructed devices like robots or either computer simulations of artificial life. Dr. Emeche still tries to situate the reader, posing a definition of Umwelt that should be immune to vitalism, but he himself doesn't escape from the prejudice in which biologists consider organisms as something that is different from artificial devices "by definition". An this should include the concept of Umwelt that, as he concludes, should apply only for real organisms or, in a very abstract sense, to a kind of artificial device to be developed in future, where some specific properties of organisms, still to be discovered should, in thesis, make us wonder that such an artificial device will be possible to exist someday.

Unfortunately, I don't share the exact same opinion. In the discussion which follows, I will try to point out the main issues in which I agree with Dr. Emeche's ideas (and this is much more than just a few points), but also will try to depict the main disagreements I have. I hope this can lead to a good and productive debate in a very controverse issue like this one being discussed here.

Umwelt and Organisms

Let's get the definition of Umwelt Dr. Emeche pointed out in his paper (actually it is T. von Uexkull's definition), and try to analyze it ... "The umwelt may be defined as the phenomenal aspect of the parts of the environment of a subject (an animal organism), that is, the parts that it selects with its species-specific sense organs according to its organization and its biological needs". Obviously, this definition actually segregates the idea of Umwelt to living organisms - those found in nature - avoiding artificial creatures to have one thing like that. But what happens if we try to adapt this definition of Umwelt in order to exchange those "biological keywords" by similar ones ? Let's experiment a redefinition of the term "Umwelt" ... "The umwelt may be defined as the phenomenal aspect of the parts of the environment surrounding a system, that is, the parts the system

is able to interface with, by means of its sensory and actuative devices, according to its internal organization and its internal objectives". Notice how this new definition of Umwelt is free from "biologisms", still retaining the same powerfulness as a concept. With this new-and-improved version of the "Umwelt" definition, we will have no problem in including artificial systems like robots as having an "Umwelt". Of course that things are not so easy like that. We have first to take account of some important issues.

First, it is tacit that organisms are really "different" from artificial mechanical devices. This difference has been widely investigated within the field of theoretical biology, especially in the work of Robert Rosen (Rosen, 1991). We will be too simplistic in saying that we will be able to fabricate an organism by means only of putting its parts together. Nobody was able to do so until now, and probably will never do. Not because of a lack of knowledge on how parts are glued together in an organism. But because organisms are really systems from a different nature. They are what Rosen calls a "complex system". One important property of "complex systems" (according to Rosen) is that they cannot be split, or decomposed, into its constituent parts, without loss of functionality. In other words, the whole is more than a set of parts working together in collaboration. Or else ... any finite model would not suffice to describe a complex system in its entirety. So, what should we do ? Well, there are actually a lot of people working on a theory on how to deal with complex systems (Mikulecky, 1995, 1999).

So, ... , what ? Well, we acknowledge that organisms are not machines ! Organisms do have something, some inner property that makes it different and not implementable by means of machine technology. But this is not the same as saying that a machine cannot have an "Umwelt" ! Let us explain a little bit our wide understanding of Umwelt.

Environments and Umwelts

Figure 1 below will help us in understanding what is this wider notion of Umwelt:
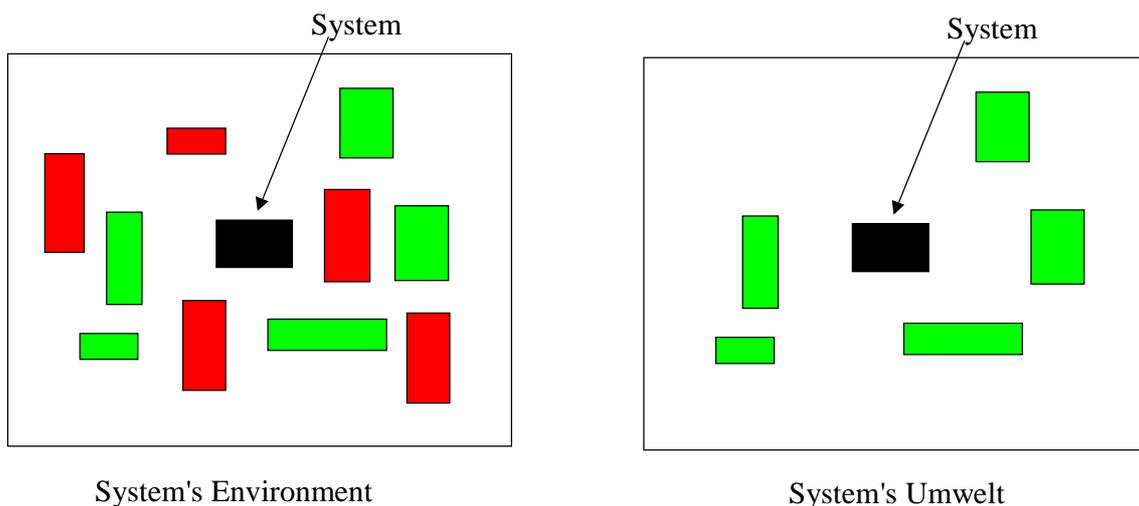

Figure 1 - Environment and Umwelt

In figure 1, we have a given system, surrounded by its environment. Suppose now that the system sensor devices are able only to capture the "green" color, i.e., it is not able

to detect red color. For this system, its "Umwelt" would be like the right box of figure 1, made only of green boxes. Notice that the "subject is NOT the constructor of its own umwelt", as Dr. Emeche said, because the Umwelt does exist in the real world, despite any influence of the system. On the other hand, we can see that there is a difference between the real environment and the umwelt, and this difference is caused by the system's sensor devices, that are able only to capture the "green" color. So, what is happening here ? The problem is that we are mixing the idea of "Umwelt" and the internal representation that the system have of its Umwelt (see figure 2) !
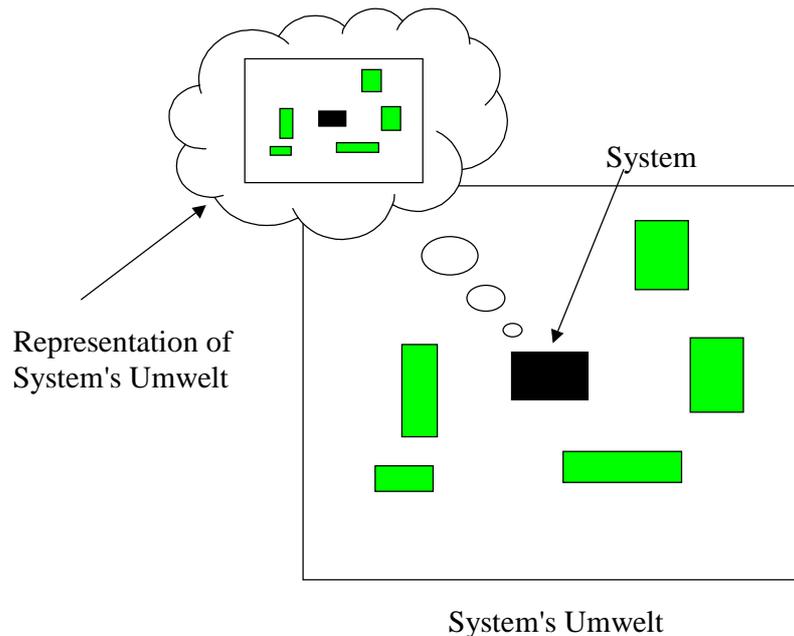

Figure 2 - Umwelt and Umwelt internal Representation

   In figure 2, we have the system, its umwelt and the internal representation that this system has of its umwelt. So, this put a little bit of light on the question whether an artificial device would have an "Umwelt". Sure it has ! The point here is that maybe it does not have an internal representation of his Umwelt !

   Let's analyze more critically those issues ... this is a little bit tricky, because of the reductionistic understanding we (humans) have of our world and our environment. Are we able to know our environment, as a whole ? NO ! The best we can know of our environment is our own Umwelt ! Is our environment made of objects with properties like color, hardness, temperature, etc ? NO ! This is only a higher-level model (internal representation) of our own Umwelt. An object that I am seeing just right there in front of me is not really there. This is only the internal representation I have for whatever is really there in the environment, that is also in my umwelt due to my ability to sense some parts of the environment. A pictorial diagram for such ideas is given in figure 3:
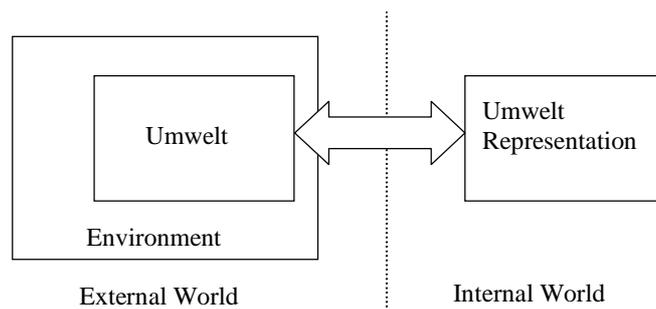
Figure 3 - External and Internal Worlds

Figure 3 shows some very interesting points. First, the existence of Umwelt is independent of the system as a semiotic agent. It is there, because it is just a part of environment. The problem now is to know what it is ! We know that it cannot be greater than the environment, because it is a part of it. But then, how to filter the parts of the environment that belongs to the umwelt and those that do not ? For that, it is critical to know which are the sensing devices of our system. This is the most important issue when determining some system's umwelt. If the system is able to know a part of environment due to their sensors, this part have to be acknowledge as being part of its umwelt. Otherwise, it does not belong to the system's umwelt.

So, it seems that the only thing necessary to characterize a system's umwelt is its sensor devices. But now, think with me ... doesn't robots have sensor devices ? So, they DO HAVE an umwelt too !

But, ... do they have also an internal representation of such umwelt ? In this case, I will have to say ... sometimes yes, ... , sometimes no !

Another important property of umwelts is that we (humans), as being systems like others, sharing the same environment, would not be able to understand parts of other organisms' umwelts, unless they also appear in our own umwelt. Suppose that some part of environment is within another organism's umwelt but is not in our own umwelt. This implies that this part of environment does not exist for us, because we are not able to feel it. Of course, this rule can be broken if we consider that men are able to "extend" their own umwelt, by means of enhancing their sensory capability with artificial sensors. This is how we are able to use and conceive electromagnetic waves in radio and TV transmission, accounting that they are not present in our "standard" umwelt.

Robots, Umwelts and Internal States

Before we can proceed, it is important to introduce some basic ideas from which we are going to derive our arguments on robot umwelts and (possibly) robot consciousness.

We will first try to get a common understanding for such words like "environment", "umwelt" and "internal representation" (people can give very different meaning for those words - some of them even contradictory). We understand the environment as taking place into a space, hereby called an external space, which can be shared with all other systems that do exist on such a place. This includes we humans and everything we humans acknowledge as taking place in our own umwelt (our version of the environment). We can always speculate that the environment is greater than simply

our umwelt, and this has serious consequences, but let's stay with the idea that our umwelt is (partially) equal to our environment, or if you wish, that our umwelt (as humans) is always greater than the umwelts of all other systems we are going to analyze, such that any system's umwelt should be within our human umwelt, that we will call "environment". Now let's focus on this external space we share with other systems. How does it look like ? Apparently, it is filled up with "something" (the things around us in the material world) that can surely be partitioned within smaller "somethings", that are understood as being "parts" of the environment. Let's remember that this illusion that things exist is only a creation of our mind. Let's then call this "something" we classify as being things in the world, as phenomena occurring in the external space, in the lack for a better word.

Now, there comes one of the trickiest parts of our reasoning. When we acknowledge the existence of some object in the world, this does not happen instantaneously. Even as organisms, we don't have "sensors of things". Sensors are only able to collect information about some quality (and this information can be a quantitative information about this quality). We can sense temperatures, lightweights, light intensity, pressure, chemical gradients, etc. We can not sense e.g. "a pencil". We can sense the contact of our body with another body, its temperature, its color. The re-organization of all such information into an abstract idea that acknowledges that "there is a pencil here" is made within our own mind. The pencil is not at the "Umwelt" it is in the internal representation of our umwelt. And this is valid both for organisms and robots.

The only information our mind is able to collect from environment (or, if you wish, our umwelt) is due to sensor information. But, ... , what are sensors ? In order to have a clear understanding of what are sensors, I will invoke Rosen's idea of "modeling relation" (Rosen, 1991 - also in Dress, 1999). Modeling relation is described in figure 4 below:
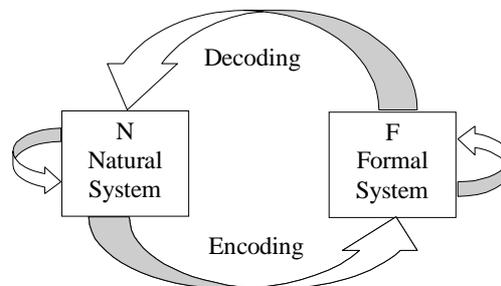

Figure 4 - Rosen's Modeling Relation

Now citing Dress' paper on Rosen's modeling relation ...

"... consider a natural system denoted by N and a corresponding formal system denoted by F. The natural system consists of components where any identification, language, or discourse leads us immediately into some formal system. N not only exists, but exhibits behaviors that can be observed by the process of measurement (Rosen, 1978). The temporal evolution taking place in such an N is an internal entailment usually identified with causality. F likewise has components in the form of elements, perhaps axioms or theorems, or maybe just embryonic concepts bearing quasi-logical

relationships to one another. The entailment structure possessed by F is formally embodied in production rules that may be identified with inference of a logical, mathematical, or algorithmic nature; F is a pure syntactical system with meaning imputed from the external context. Entailment is an internal, syntactical operation within F that, in the best models, is an analog of the autonomous, causal, and dynamic processes in N. The measurement process is subsumed by the concept of encoding behaviors of N into elements or sets in F, often in the form of parameters. Inferences are made in F by 'turning' the mathematical 'crank.' The results of the inferences must then be compared to the future behavior of N, which process is called 'prediction' in science. Decoding is a dual process to encoding—one that the scientist uses to verify predictions by another measurement procedure on the future behavior of N."

Despite Rosen's modeling relation is conceived, in principle, to deal with the comparison between a natural system (the environment) and a formal system (the internal representation of environment), this idea can be used in order to describe and understand the nature of umwelt.

Rosen say that if we are able to set up an encoding function that for elements in N assigns respective elements in F, and also a decoding function that for elements from F assigns respective elements in N, then we build a modeling relation among N and F, in the sense that F models N in respect to the parts of N and F that are used as domain/counter-domain for encoding and decoding functions. If we are able to commute the encoding and decoding functions (exchange one for the other), than we are able to say that F is an "exact model" of N. Otherwise, it is simply a "simulation" of N with respect to the encoding and decoding functions being employed.

This idea is very interesting, because the most basic types of encoding/decoding that can be employed will be implemented through sensor devices. This leads to a property in which for each "part" of environment (N) that can be "sensed" by a sensor (encoded), there should be a corresponding part in F that is modeling that part in N. So, if we have such modeling relation that is set-up using sensor/actuators as encoders/decoders, then there should be an element in F for each part in N that can be sensed/actuated. This sets up a system F that is a *simile*, or in other words, an "exact model", not of the environment, but of the system's umwelt ! This because the umwelt comprises only the parts of environment the system is able to sense, and if the system is able to sense, there should be an encoding/decoding of it from N to F. So, sensors and actuators are the perfect "filters" to be used when going from environment to umwelt. More than this ... we know that, in thesis, it is possible to grow up an "exact model" F of a system's umwelt N.

Now, let's try to understand what our system is able to do with this information it gets from umwelt. It can do basically three things with the information sensed.

- It can use this information to create (and update and store) a **model of its environment** (actually, a model of its umwelt), OR
- It can use this information to build some sort of an **evaluation**, that is going to be used to judge if the **system objectives** are being accomplished or not, OR

- It can use this information to generate a **behavior** to be performed by the system.

Different organisms and artificial devices would perform these functions in different ways, sometimes not implementing the whole three.

Let's analyze first the model creation functionality. What is this environment model that is created and updated by the system ? Actually, it is an icon of the system's umwelt. But more than a simple icon, this icon can be a perfect *simile*, an exact model, of the umwelt. But, ... , do all the world models are really *similes* of their umwelt ? No ! More than this ... do all systems really STORE their world models ? No, again !

This is the source of the whole polemics on robots and umwelts. Because Dr. Emeche doesn't differentiate the real umwelt from the internal representation a system makes for its umwelt, he mistakenly concludes that a robot may NOT have an umwelt ! If we make this differentiation, we can clearly see that all robots do have an umwelt, despite the internal representation carried out by a particular robot may not adequately represent this umwelt. Or either, the internal representation used for its umwelt does not comprise a "memory" of it.

Some autonomous agents (mainly the called reactive agents, like those pointed out by Dr. Emeche - studied by Brooks and others) actually do not carry an internal representation (a memory) of their umwelt. They use their sensed information to instantaneously generate a behavior, which is performed by the system. But we showed (Gudwin, 1996), that for some kinds of problems (and even simple problems), the reactive paradigm can not leads us too far, without some kind of planning. And planning is only possible when we have a model of our umwelt inside the system.

This sets up our attention to another important issue regarding artificial semiotic processing (sometimes provocatively called ... robot thinking): how do they store their umwelt model ?

Organizing Knowledge

In building artificial semiotic devices, one important constraint is the available physical memory in order to store the system's umwelt. As the environment is (hypothetically) infinite, even with a finite number of sensors, those sensors would be located in a potentially infinite number of points within the environment. That would require an infinite memory in order to store all the system umwelt. This is neither efficient nor feasible. If we want to create an artificial semiotic device (a robot), we will need a more efficient way of storing the world model. This is when we recover the "object" metaphor we introduced some sections later. Instead keeping the information of all positions sensed by our device, paired with the position it was sensed and the time it was sensed, our intelligent system uses the "object" metaphor. This metaphor says that if we are having a kind of sensing that is within some kind of periodic or repetitive pattern, either in space or in time, then maybe it occurs due to the fact that there is such a thing as an "object", with a set of attributes, that is the responsible for the repetition in the sensed signal. So, instead of storing all the sensed signals in the system history, we substitute it for a model of such "object", which the system presumes to exist in environment. The creation (or discovery) of objects can be enhanced with the positivistic/reductionist idea that a whole can be broken up in simpler parts, and so an object may have parts that are

also object, but with simpler attributes. This process can be applied over and over, until we finish with very simple descriptions of our elementary objects, which can then be put again in a model relation with N in a fairly good approximation. This will constitute a good model for planning and predicting the behavior of the environment.

Objectives

In order to be an autonomous agent, though, a system may not only have a good model of its umwelt. It must have objectives. There should be some type of goal driving the behavior of such autonomous agent. Such goal can be as simple as "live and make copies of yourself", or more sophisticated as performing some service for humans. So, the identification of system goals is one important perspective when we are building artificial devices. And our autonomous robot should be processing this kind of information as much as the information from umwelt it collects from sensors. More than this, ... , there is a closer association of umwelt and objectives, because normally we have to describe our objectives using the umwelt as a point of reference. So we have to evaluate parts of the information a system gets from his umwelt, in order to transform this information into evaluations of the umwelt ... "this is good" ... "this not so much", and so and so. For natural organisms, those objectives are implicitly encoded into organisms by natural selection and evolution. But for artificial devices, they should be conceived and developed by the human that is designing the system. It is the "a priori" knowledge that is a requirement for the system to work properly.

Clearly, objectives go beyond the umwelt, despite requiring the umwelt as a ground on which to be defined.

Behavior

Again, in order to become an autonomous agent, it is not enough to have a model of the umwelt and a mapping from and to the umwelt and the systems objectives. The agent must act, it must perform some kind of behavior in the system in order to become a truly autonomous agent. The simplest behavior is the reactive behavior, in which an information signal sensed from environment is transformed into an actuation signal and from there to an action over the environment. Despite Brooks and others have studied extensively reactive behavior, there are clearly limitations in the type of intelligence we are able to achieve purely with this kind of behavior. For more sophisticated problems and situations, there should be some type of planning, i.e., the investigation of alternative possible behaviors over predictions on how the environment is going to be after a set of possible actions. This is the main issue behind artificial intelligence and forms of emulating plausible reasoning techniques.

The Future of Robotics and Intelligent Systems

It is very difficult to preview the future, making significant assertions and at the same time accurate ones. But we are able to feel some tendencies and point out how they can be connected to building artificial devices (robots) that are more alike to real organisms, in terms of autonomy and intelligence.

One point that is very much clear is that there is a wide "gap" between "grounded" robotics, (i.e., robotic behavior that is closer to the environment through the use of sensors and pattern recognition techniques), and symbolic computations performed by

traditional artificial intelligence. Only purely reactive robotics or pure symbolic computation machines will not be able to develop the new generation of intelligent systems that will be serving us (I hope !) in the far (near ?) future. The creation of a new generation of intelligent autonomous systems will have to be inspired somewhat in semiotics. We will have to develop true "semiotic systems", and to equip them with the whole arsenal of sign-processing we are able to detect in highly evolved organisms like man and other animals. But more than this ... we will have to create a new semiotics, that is able to perform not only "semiotic analysis", but also a "semiotic synthesis". When we are studying real organisms, it is easy to give names to structures and patterns we discover, but sometimes it is not that easy when we are building, designing such artificial semiotic systems. Semiotics must go beyond a pure analysis artifact to become a true palette of design, in order to perform synthesis of new systems. We have to study not only the types of signs and how to classify them, but how they are generated and how signs are transformed from icons to indexes and from indexes to symbols and from symbols back to icons. How do we pass from signals to true signs. How those things do occur in sequence and in parallel. How concepts like the "Umwelt", deduction, induction and abduction can be useful in order to improve our robotic architecture. How are we able to build true organisms, out from complexity theory. How are we able to fabricate life ! These are the challenges for the future. These are the questions we have to solve, in order to be able to create new kinds of machines that are really more "intelligent", and maybe could help humanity to create a better world to live in.

References

(Dress, W.B. 1999) - "Epistemology and Rosen's Modeling Relation" - http://hypernews.ngdc.noaa.gov/hnxtra/Dress_Paper.html

(Emeche, C. 1999) - "Does a robot have an Umwelt ? Reflections on the Qualitative Biosemiotics of Jakob von Uexkull" - 2o. Seminário Avançado de Comunicação e Semiótica - Novos modelos de Representação : Vida Artificial e Inteligência Artificial - São Paulo, 18-20 agosto de 1999.

(Gudwin, R. 1996) - Contribuições ao Estudo Matemático de Sistemas Inteligentes - Tese de Doutorado - DCA-FEEC-UNICAMP, Maio de 1996.

(Mikulecky, D. 1995) - "Life, Complexity and the Edge of Chaos : Cognitive Aspects of Communication Between Cells and Other Components of Living Systems" - submitted to Biotheoretica ACTA, Nov. 27, 1995.

(Mikulecky, D. 1999) - VCU Complexity Research Group Home Page - http://views.vcu.edu/complex/

(Rosen, R. 1978) - "Fundamentals of Measurement and Representation of Natural Systems", Elsevier, 1978.

(Rosen, R. 1991) - "Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life" - Columbia University Press, NY, 1991.