A Biosemiotic Framework for Artificial Autonomous Sign Users

Erich Prem

Austrian Research Institute for Artificial Intelligence
Freyung 6/2/2
A-1010 Vienna, Austria
erich@ai.univie.ac.at

Abstract

In this paper we critically analyse fundamental assumptions underlying approaches to symbol anchoring and symbol grounding. A conceptual framework inspired by biosemiotics is developed for the study of signs in autonomous artificial sign users. Our theory of reference uses an ethological analysis of animal-environment interaction. We first discuss semiotics with respect to the meaning of signals taken up from the environment of an autonomous agent. We then show how semantic issues arise in a similar way in the study of adaptive artificial sign users. Anticipation and adaptation play the important role of defining purpose which is a neces-sary concept in the semiotics of learning robots. The proposed focus on sign acts leads to a se-mantics in which meaning and reference are based on the anticipated outcome of sign-based interaction. It is argued that such a novel account of semantics based on indicative acts of refer-ence is compatible with merely indicative approaches in more conventional semiotic frame-works such as symbol anchoring approaches in robotics.

Introduction

Issues of semantics have a long history in the study of adaptive and evolving systems. Ever since the seminal work of Uexküll (1928) in biology researchers were interested in the explanation of how something like "meaning" is created in system-environment interaction. In our days, modern system scientists (e.g. Pattee 86) addressed these questions. In the area of Artificial Life (ALife) it is often robotics researchers who focus on problems related to signs and artificial systems: work in this field ranges from stigmergic communication (Mataric 95) to behaviour-based robots using signs and language (Steels & Vogt 97, Steels 01, Billard & Dautenhahn 97, 99). In particular, technical approaches to mapping objects in an autonomous robot's environment on structures internal to the robot ("Symbol Anchoring") are an active field of research (Coradeschi & Saffiotti 03).

It is all the more surprising that recent publications in this area rarely address foundational is-sues concerning the semantics of system-environment interactions or other problems related to biosemiotics (a notable exception is (Cangelosi 01)). As we will discuss below, the approaches make little or no reference to specifically life-like or even robotic characteristics such as goal-directedness, purposiveness, or the dynamics of system-environment interaction. These features are, however, central to much of robotics and ALife. It is highly questionable, however, whether technical approaches to symbol anchoring should be developed devoid of any sound theoretical foundation for concepts such as "meaning" or "reference". Until now, simplistic versions of Fregean or Peircean semiotics seem to have motivated existing technical symbol anchoring and symbol grounding proposals.

This is all the more regrettable, since robotics lends itself nicely as a tool for the study of se-mantic processes in life-like systems. Robot and ALife models offer the potential for systematic in-depth analysis of complex system-environment interactions where many (or all) parameters are known, simply because these systems have been constructed by their designers. For this approach to develop its full potential, however, it is necessary to first get a thorough understanding of the phenomenon under study. This is the aim of the work described here. It is performed in the con-text of learning robots in which the acquisition of object concepts and learning names for these objects plays a central role.

In this paper, we investigate the nature of signs and semantics in relation to robotic systems. In particular, we propose a framework for the study of semantic aspects of autonomous (artificial) sign users with the aim of clarifying concepts such as "reference" and "meaning". We investigate the role symbols and other signs play in autonomous systems and address biosemantics from an ethological perspective so as to develop a system theoretic framework which also reconciles symbol anchoring with sign-act perspective of meaning in robotic agents.

Semiotics in ALife

Semiotics generally refers to the study of signs. In a psychological (and semiotic) context, signs usually are physical objects that carry meaning for humans. From a more biology-oriented per-spective, signs are simply signals that carry meaning. In what follows, we will use both characteri-sations based on the assumption that physical signs also generate “signals” that can carry meaning in this sense.

The referent in traditional semiotics is what a sign stands for, e.g. “cow” stands for a cow in many contexts and the same is true for the German word “Kuh”. While syntax studies the relation of signs to other signs, semiotics deals with the relation of signs to their “meaning” or, more pre-cisely, their referent. Pragmatics then covers all the aspects that relate to the actual use of the sign by its interpreter. In what follows we will slightly blur this seemingly clear distinction between pragmatics and semantics. This is a natural consequence of focusing on life-like systems that ex-hibit purposeful behaviour in their environment and either actively use or passively “read” signs.

In the context of robotic, ALife, and AI systems, the interest in signs and their meaning arises from at least three different perspectives. The first originates in the aim of creating a system that uses signs for communicative acts with humans or other artificial systems. The underlying motiva-tion here can either be to achieve this desired communication or to study processes of child-like language acquisition or even, language evolution. The second perspective – usually found in ro-botics – is to connect the meaning of internal structures (“representations” or “anchors”) to objects in the world. Here, the goal often is to create a model of the robot’s environment for planning purposes. Finally, another perspective (often alluded to by the two others) focuses on the more philosophical “Symbol Grounding Problem” that arose in discussions following John Searle‘s famous “Chinese Room” argument. Harnad poses the question as to how it is possible for an arti-ficial agent to acquire symbols that possess intrinsic meaning (Harnad 90, 93) which is not “para-sitic” on the meaning of other symbols in our head.

In many ALife approaches, all these efforts result in the challenge of how to establish and main-tain a relationship of sensory data of objects in the agent’s environment with symbolic representations. The result of such a process are descriptors for sensory signals that allow the classification of a part of the sensory stream as caused by an object in the environment of the agent. In the case of “anchoring” the focus lies on proper technical solutions ranging from prototypes, feature-based approaches, to more sophisticated dynamical systems solutions (cf. Davidsson 95). In the case of making an agent use and “understand” human language, the focus is on the relationship of the agent’s categories and human word use (or “word use” of other robots).

In both cases – and this is the relationship to symbol grounding – the assumption usually is that there exists an internal mediating representation that captures the “meaning” of symbols (either as words or as program-internal constructs). In most examples described in the literature, symbols refer to static objects or to features of these objects. Verbs and other word categories are only rarely grounded in practice.

These approaches often root in a naïve view of language as a means of communicating informa-tion about situations to other agents, as if the purpose of language would be to inform about a state-of-affairs. In Alife models the aim then is to automate the construction of such a language by an autonomous adaptive agent. In essence, this amounts to the automated generation of models of the environment that happen to be understandable by e.g. humans, cf. (Prem 00). The semantics of the sign tokens used thus are components in such an image-oriented model in which the sole pur-pose of signs is to “represent” object-like states in a model of the environment. In what follows we propose a completely different approach to semantics for autonomous artificial sign users.

Active Sign Users

In contrast to the image-oriented view described above, we will focus on the semantic action (and thus include aspects of what is usually termed pragmatics) in the analysis of artificial sign users. Following the approach of the philosopher Martin Heidegger (“Being and Time”), the semiotician Charles W. Morris (“Signs, Language, and Behaviour”), and the work of the linguist J.L. Austin (“Doing things with words”), we regard signs as tools that enable agents to pursue their goals. Consider the following examples of using signs or “sign acts”:

Sign Act	Behaviour
Greeting	The agent reacts to a greeting or salutation or to another agent with a specific greeting behaviour.
Set mark	The agent marks an interesting location or object in the environment so as to retrieve it later more easily.
Warn	Produce an alarm signal to make group members aware of danger or make them run away.
Flee	React to a warning signal by running away.
Follow arrow	The agent moves in the direction to which an arrow points.
Find place	The agent navigates to a designated place. Examples include “here”, “there”, “home”, etc.

Table 1. Examples of sign using behaviours (sign acts).

This list can be easily extended to include more language-like behaviours, e.g. “A beer, please!” to make a waiter carry beer to your place. These sign acts are understandable because their *purpose* is immediately clear to us human observers. And for systems that pursue completely different goals, they would probably not make much sense. This idea has been nicely captured in Wittgenstein’s famous word “If a lion could speak, we would not understand it.” (Wittgenstein, 1953)

Note that the “purpose” here refers to the outcome of the sign act and lastly is a feature of the system-environment interaction. In a less teleological terminology, the purpose simply is a consequence of a selector operating on a “sequence” of similar systems so that there exists a function ϕ, which the system can be said to minimize. In ALife, ϕ is often called the “fitness” of the individual.

Indicative Meaning

The proposed focus on sign acts leads to a semantics in which meaning is defined as the anticipated outcome of sign-based interaction. The idea here is that an agent uses a sign because of the anticipated outcome of the indication action. Heidegger’s example is a driver who uses the turning signal of a car to indicate the intended turn. The assumed intention behind this sign act is based on the driver’s anticipation of less problematic driving due to the reaction of other drivers. On the other hand, the other car drivers use the turning signal appropriately and guide their vehicles so as to avoid collisions etc. It can now be argued that the turning signal thus refers to a whole network of activities in the context of driving cars. The purpose of the sign is to give orientation to these networked behaviours of the sign-users.

Note that this example of the car turning signal is a challenge for a simple semantic approach that focuses on mapping external entities onto signs. It is not at all clear to which “external entity” a turning signal could be said to refer. The remainder of this paper serves to clarify this intuitive argument for a semantic approach that focuses on indicated outcomes of sign-based interaction. We consider the framework to be in line with what has recently been termed interactionism (Bickhard & Terveen 95 and Bickhard 98).

Autonomous Sign Users

Autonomous Adapted Systems

Ethology provides a set of useful tools and constructs for the analysis of autonomous systems, such as e.g. robots. As early as 1928 the theoretical biologist Jakob von Uexküll proposed to analyse living beings following a close study of the interaction of the system with its environment and of a detailed analysis of their sensors and actuators.

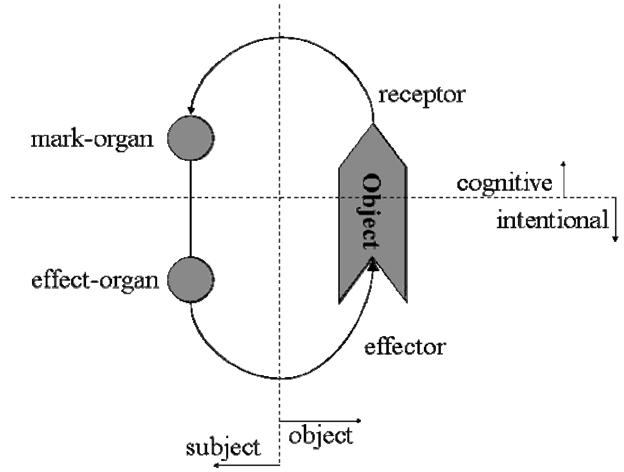

Fig. 1. An action circuit: the properties of the mark organ and of the system-environment interaction determine the world of the animal. See text for further description.

Figure 1 illustrates the action circuit as an important construct in von Uexküll’s work. It depicts objects in the agent’s environment as perceived and interacted with by the agent. This perception is determined by the sensory system of the animal (the “mark organ”) and the interaction is determined by the effectors. The important point here is that the object according to the agent is nothing but the marks perceived and the potential for interacting with it. A frequently used example is the (simplified) tick which feeds on warm-blooded animals. The tick simply releases its grip when it perceives a significant concentration of butyric acid. It then bites objects with a temperature of 37 degrees centigrade. Thus, there are no animals in the world according to (hungry female) ticks, only butyric acid emitting warm possibilities to seize and bite. The circuit as depicted here also nicely illustrates the dichotomies between subject and object and between the cognitive and intentional aspects of the interaction.

Behaviour-based robots can be described using the concept of the “action circuit”. The behaviour-based architecture (Brooks 91) consists of a set of layered modules. Each of these interacts with the environment and is capable of driving the robot. The modules all read from sensors that are highly tuned to the specific behaviours and their output can either directly influence the robot’s behaviour or indirectly suppress other behaviours. As a consequence, behaviour-based robots interact with the world at high interaction dynamics. The similarity of this approach to early ethology has been described before (e.g. by Prem 98) and can be exploited when analysing robots and – more efficiently – when designing the sensor and effector systems of a new robot for some specific task.

A Biosemantic Analysis

Following the description of Bischof (95) the environmental interaction circuit lends itself nicely to a semantic analysis in system theoretic terms. "Semantics" here of course refers to the conditions or rules under which it is justified to say a sign refers to a given state-of-affairs (or, more complex, whether a given sentence is true). Morris argues that

> if something, A, controls the behaviour towards a goal in a way or similar to (but not necessarily identical with) the way in which something else, B, would control behavior with respect to that goal in a situation in which it were observed, then A is a sign. (Morris 1946)

In Uexküll's example of the female tick there are two important pathways in a simplified system theoretic description of its behaviour. On the one hand, the chemical concentration of butyric acid is a measure of the distance of the animal to the tick. As soon as this concentration is above the threshold, the tick drops and bites. On the other hand, the distance from tick to warm-blooded animal influences the probability that a dropping tick actually hits its victim. This probability is greater than zero only within a narrow range that corresponds to the animal being close to the tick. The connection between these two pathways (the internal release mechanism and the "external" property) defines the expected probability for a successful feeding interaction and thus influences the probability for this animal to survive.

In this way we discover ("semanticize") that chemical concentration is a stand-in for the warm-blooded animal. Note that the reference to *purpose* is essential in this context. As mentioned before, "purpose" really refers to the function implicitly created by the selector that operates on a series of similar systems.

If you prefer a more technical example, consider the case of a vacuum-cleaning robot homing in on its battery charging station and an adequate release mechanism such as making contact with the charger. A candidate for the "purpose" would be the selector function that has in the past lead to a generation of robots which successfully charge and are thus bought by customers.

In what follows, we use Bischof's more formal framework. It depicts the interaction using Mason diagrams in which signals are represented as circles and transfer elements as arrows. For Uexküll's original example of the tick, the semantization looks as follows (see Figure 2).

We depict the relationship between the environment (α) and the action of the agent (ω) selected to improve on fitness (ϕ). For the sake of clarity and to make the relation to Uexküll's original picture clear consider that signals within the agent are separated into (s) at the sensor side and (r) closer to the effector. The signal r can take two different values ("drop") and ("stay"). The signal s corresponds to the perception of butyric acid, in principle using a continuous range of values. For our purpose here, however, only "strong" and "weak" are interesting cases. (The meaning of α' is described in the next section.)

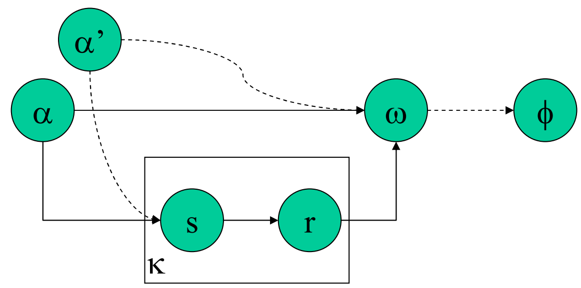

Fig. 2. A Mason diagram depicting original and sign-generated signal flow in Uexküll's tick example – adapted from Bischof (95). The autonomous system κ perceives signals via s from its environment (α) and interacts through r so as to generate changes in the environment (ω). The connection from α to ω denotes a homeostatic relation, which in turn influences the selector ϕ. α' depicts the sign. See text for further description.

The environment of course can be described in many ways, but taking into account the action circuit and the possible ways of interaction, only two interesting descriptions remain for (α): "warm-blooded animal & close" and "no animal or not close enough". We now have four interesting system-environment interactions that are a result of the possible values of α and r: a matrix for (ω) with the values (first row) "feeding" and "rest" and (second row) "miss opportunity" and "waste energy".

The cognitive meaning of "strong" then is to look for the corresponding value of ω in the first row and search for the α which maximizes ϕ for the possible values of s. The result is "warm-blooded animal & close".

A similar argument gives the intentional meaning "feeding" for "drop". Note that this is already a semantic interpretation of the agent's interaction that takes into account the environment, the agent's interaction properties, and a "biological" fitness function.

There are two interesting consequences of this framework. First, it is now possible to give clear meaning to discussions about deception or error (see Bischof 95). More importantly, however, we can now take a more precise look at signs and sign acts.

Towards a Semantics of Autonomous Sign Usage

The tick does, of course, not use physical signs in the more conventional sense of the world – only signals. Nevertheless, it should be pointed out that the analysis of Uexküll and of many biosemioticians interprets the sensory signals as signs similar to what we have outlined here. There are, however, robotic examples of sign users similar to this example.

Concentrating on the cognitive part of Figure 1, let us look at a robot which perceives a specific sign, such as foot-waving in the famous lab-guide MIT robot "Polly" (Horswill 93). On perceiving a waved foot, the robot starts giving a tour to visitors. It should be straightforward that in this case the analysis of the system will follow precisely the same steps as in the tick example. The situation is really not much different and this is the reason why a separation of the physical sign and its "signal" does not add much to an understanding of the meaning of the sign.

Following Morris, it is now possible to clarify better what it means for a sign to "stand in" for an object: Take as another simple example the flag that marks the hole in the putting green of a golf court. The flag points out the location of the hole and thus (hopefully) makes players orient their stroke towards this sign. In our terms, the sign gives orientation to the interaction as it stands-in for the hole (in the context of playing golf). *The sign gives orientation to an action circuit.*

More formally, we can simply take the situation depicted in Figure 2. Here, α' controls behaviour towards the goal (ϕ) similar to the way in which α would control behaviour with respect to the goal in a situation in which it were observed. This, however, is exactly Morris' definition of a sign.

As a side-remark notice that this analysis also clarifies the difference between the three types of Peircean signs, i.e. icons, indices, and symbols. In the case of the *index*, there is some "inherent" connection between the sign and its referent. Peirce uses the example of the arrow on top of a roof that is a sign for the direction of the wind. In Figure 2, an index α' will have a law-like relation with α external to the agent. The index may be recognised using a completely different s, but the connection of α and α' will be based on a causal connection. For an *icon*, i.e. an image-like similarity, the connection of sign and referent will be based on characteristics internal to the agent, i.e. on the features of s when perceiving the sign. The characteristic feature of a *symbol* is the arbitrary connection that is only connected to the system-environment interaction and its results, i.e. to the optimization function ϕ.

This analysis allows the formulation of a research program to clarify the meaning of signs encountered by autonomous agents. The program consists in discovering the relation of the entities depicted in Figure 2, their influence on "fitness" and the way in which signs stand-in for already semanticized signals. An example is the man who we observe looking for something in an unknown place. When we see him entering a room that is marked with "Cualquatzi" we may decide after some more observations that this word stands for the men's room. Following the same line of arguments as before such a deduction would be based on the relationship of the estimated purpose of the men's interaction (ϕ) with its environment when looking around.

Adaptive Anticipatory Sign Users

Our discussion has so far focused on passive sign usage. We shall now turn to signs that are actively created by autonomous agents and put them in the picture which we have used so far.

Active Sign Users

A first example for an active, autonomous sign using robot would be a machine that cries for "Help!" to get out of places where the robot is stuck. This would be a usage of the sign that can be perfectly described using an action circuit. The action would be the utterance of the sign "Help" and the question is what the meaning of "Help" is. In more semiotic terms, this question amounts to the conditions under which the utterance of "Help" is "true" (let us, for the time being, say "useful"). As proposed in the introduction and using imprecise language here, the meaning of "Help!" is the anticipation of getting out of the current problematic situation. (Also, consider the case, where the one-word sentence "Mom!" is used by babies to receive attention and help and *not*, of course, to *"refer to"* or *"label"* mothers.)

In the formalism used in Figure 2 only the values for r and ω need to be changed. The agent's action r now comprises sign acts, e.g. a linguistic "action" of word utterance. The rest of the action circuit stays the same, in principle. In such a view, language and other sign acts are simply regarded as behavioural interactions with the environment. The question arises then, how it is possible to pin down the meaning of the signs used in the "intentional" part of the action circuit as depicted in Figure 1.

A possible answer to this question is to simply investigate what the sign means for other autonomous systems that use the sign in their action circuits. This, of course, amounts to searching for those values of α which are in Morris' sense similar to α' for agents that take the sign as input. Their understanding of the sign, however, may or may not be the same as what was originally meant. Another answer is to simply use the approach taken in section 2.2 and search for the maximum value of ω which corresponds to the sign action r.

Apparently, this notion of meaning for the use of signs such as "Help!" is not very intuitive. Indeed, we encounter the same situation as in the case of the turning signal where the meaning is not easily describable in object-like terms. The reason for this difficulty lies in the fact that we are using a specific purpose of the interaction here. Using the turning signal amounts to anticipating successful outcomes of indicating one's direction to others, i.e. to influence the behaviour of other participants in traffic appropriately. These participants in turn also use the signal in orienting themselves towards the sign.

The action circuit of "driving home" etc. receives its orientation through the sign that is actively used by the driver in front. Our discussion of reference in autonomous sign users is based on the tool-character of signs and not on pure reference. Reference in the discussion here is not something that happens merely because signs are used to refer. In contrast, signs are used for circumstances and purposes in which (we use a terminology close to Heidegger (27) and Dreyfus (90) here) an agent already dwells. Signs are merely another tool to achieve the desired interaction outcome.

Signs are, however, peculiar in pointing out parts of the contextual whole of references in which an agent might already find itself. What Heidegger calls the "wherein" of living can be analysed as the action circuit(s) which receive(s) orientation using the sign.

We thus propose to regard active sign acts, i.e. acts of creating or putting signs as anticipations of successful interactions of indication. Passive sign usage, i.e. following or taking up signs encountered in the environment on the other hand should be regarded as giving orientation to action circuits. Both views are compatible with a biosemiotic analysis which allows discovering precise meaning even for signs which are mainly used as environmental actions and not for "pure" reference. In section 3.3 below we will reconcile "pure" reference with this view.

Adaptive Autonomous Sign Users

Consider now the case where an autonomous sign user is adaptive, e.g. trained using a machine learning algorithm or artificial neural network. We know that in such systems there usually is an error function which is minimized based on some feedback from the environment (or a teacher). In effect, this feedback is based on the system's previous actions. Thus, what happens in adaptive systems is that the system generates an implicit model of the environment and its interaction with it. In the case of a neural network for example, the network weights will after a while represent this model.

It is important to realise that there are at least two different time scales present in such a process: on one level, the system received feedback from the environment based on the selection of actions. At another level parameters of the model that generates the actions are adapted. The latter process necessarily operates at a slower time-scale. After some time, the effect of this selection mechanism operating on the system parameters will yield a model that generates actions which appear to minimize future error, i.e. the system takes actions now which will generate the desired effect later. This simple effect follows from the adaptivity (and to some extent from the autonomy) of the system.

In our terminology, the system will adapt with respect to a function that also minimizes ϕ. After some period of adaptation, it will appear to select actions so as to minimize ϕ. Assuming that the agent adapts its internal structure (s and r or whatever this structure looks like), an appropriate selector function will thus yield systems which "decide" to use signs *now* so that the desired outcome of the sign-act follows *later*. It thus follows that signs in an adaptive autonomous system using signs are anticipations of successful outcomes of indication actions. The meaning of these signs needs to be explained with respect to the selector function and the agent's situation. In the turning signal example, the meaning needs to refer to the whole network of interacting agents driving their cars.

Back to Simple Reference

Let us finally take another look at reference to objects in sign acts. A typical case would be a robot that asks for a specific object or indicates (points out) an object to another agent. In these interactions with the robot's environment, the anticipated outcome of the interaction can mean a successful indication of an object.

For example, when asking "give me the soda can on the table" the robot refers to the soda can in the sense that it anticipates to be given the soda can successfully. The term "soda can" in this context refers to the soda can because a successful interaction means that in this case the soda can is given to the robot. The point of this seemingly trivial circumscription is, of course, that for certain acts of interaction with the environment, it is possible to discover a simpler correspondence of aspects of the environment (□) and the sign based interaction r. This is the case for those interaction circuits which are primarily indicative in the traditional sense.

The question then arises how such purely indicative signs could be created in artificial adaptive autonomous sign users. We still will have to start with more goal-driven sign acts. Let us, for example, assume that the first "descriptive" nouns (something with a meaning like "mummy" or "daddy") in baby talk are produced so as to generate parental reward, *not* to denote the corresponding "objects". It could then be possible in an ALife model to devise a method for re-using a representation anticipating this kind of reward for other rewards or purposes (ϕ).

As an example, it would be interesting as a next step to also receive reward from others for correctly pointing to mum and dad or for crying for help etc. The essence of all these potential anticipations of internal indication outcomes, could then converge towards a mainly referential representation that happens to denote mum and dad in a number of contexts.

Staying with baby or robot talk for a moment, note that there is a close relation of “Give me X!” and “This is X!”, because the former sentence happens to produce the desired outcome, if “X” refers to X. Note, however, that “X” can still mean many different things in other contexts. Similarly with vervet monkeys, “leopard!” will model the leopard, but always in order to make one’s friends jump up the trees. As a consequence, it is not necessary to enrich “connotations” of the word “leopard” after its descriptive properties are learned. They will, quite to the contrary of such a view, ensure that the sign is properly used in the first place.

The problem with many existing approaches to the study of signs in robotic models is that they start with “This is X!” as the *only* indicative interaction of the autonomous agents. It is obvious, that such an approach does not leave any room for the more purpose-driven adaptive sign acts discussed here.

Summary

In this paper, we have described a biosemiotic approach to the analysis of autonomous sign users. We have outlined a framework for studying the semantics of signals and different kinds of signs using simple system-theoretic concepts. It was argued that adaptivity in such systems necessarily leads to anticipatory sign usage. Finally we have shown how naïve reference still fits in this picture when talking about a specific subset of indication actions. It is this restriction to this latter subset of sign acts which so heavily influenced scientists in ALife.

It is obvious that this framework only marks a starting point for further analysing semiotic processes of autonomous systems in ALife. Developed properly, however, we hope that this framework can assist in disentangling the conceptual confusion in ALife research when it comes to topics such as language, signs, and symbol grounding.

Acknowledgements

This research was supported by the European Commission’s Information Society Technology Programme project “SIGNAL – Systemic intelligence for artefacts that live and grow” IST-2000-29225.
Partners in this project are University of Bonn (Germany), Napier University (United Kingdom), Istituto di Elettronica e di Ingegneria dell’ Informazione e delle Telecomunicazioni at Genoa (Italy), and the Austrian Research Institute for Artificial Intelligence which is also supported by the Austrian Federal Ministry for Education, Science, and Culture.

References

Austin, J.L. 1975. *How to Do Things With Words.* Cambridge: Harvard University Press.

Bickhard, M. H. and Terveen, L. 1995. *Foundational issues in AI and Cognitive Science,* Elsevier.

Bickhard, M. H., 1998. Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(2), 179-215.

Billard, A., Dautenhahn, K. 1997. Grounding communication in situated, social robots, *Proc. of TIMR 97*, Towards Intelligent Mobile Robots Conference, Tech. Rep. UMCS-97-9-1, Dept of Computer Science, Manchester University.

Billard, A., Dautenhahn, K. 1999. Experiments in learning by imitation – grounding and use of communication in robotic agents. *Adaptive Behavior*, 7 (3/4).

Bischof, N. 1995. *Struktur und Bedeutung* (Structure and Meaning), Huber, Bern.

Brooks, R.A. 1991. Intelligence without representation. Foundations of Artificial Intelligence, *Artificial Intelligence*, 47 (1-3).

Cangelosi, A. 2001. Evolution of Communication and Language Using Signals, Symbols, and Words, *IEEE Transactions on Evolutionary Computation*, Vol.5 (2), pp. 93-101.

Coradeschi, S., Saffioti, A. 2003. Perceptual anchoring – Anchoring symbols to sensor data in single and multiple robot systems, *Special Issue of Robotics and Autonomous Systems*, Vol. 43 (1-2).

Davidsson P. 1995. On the concept of concept in the context of autonomous agents, *Proc. of the 2nd World Conference on the Fundamentals of Artificial Intelligence (WOCFAI-95),* pp. 85-96.

Harnad, S. 1990.The symbol grounding problem, Ph*ysica D*, 42, pp. 335-346.

Harnad, S. 1993. Symbol grounding is an empirical problem, *Proc.of the 15th Annual Conference of the Cognitive Science Society*, Boulder, CO, Lawrence Erlbaum Associates, pp. 169-
174.

Heidegger, M. 1993. *Sein und Zeit. (Being and Time.)* Tuebingen: Niemeyer.

Horswill, I. 1993. Polly: A vision-based artificial agent. In *Prodeedings of the AAAI-93*, Washington D.C..

Mataric, M. 1995. Issues and approaches in the design of collective autonomous agents, Robotics and Autonomous Systems, vol. 16, pp. 321--331, December.

Morris, C.W. 1946. *Signs, language, and behaviour*, New York, Prentice Hall.

Pattee, H.H. 1986. Universal principles of measurement and language functions in evolving systems. In J.L.Casti, A.Karlqvist (Eds.), *Complexity, Language, and Life: Mathematical Approaches*, Springer, Berlin, 1986.

Prem, E. 1998. Semiosis in embodied autonomous systems, *Proc. of the IEEE International Symposium on Intelligent Control*, IEEE, Piscataway, NJ, pp.724-729.

Prem, E. 2000. Changes of representational AI concepts induced by embodied autonomy. In: Communication and Cognition – AI, the journal for the integrated study of AI, *Cognitive Science and Applied Epistemology*, Vol. 17 (3-4), pp. 189-208.

Steels, L. 2001. Language games for autonomous robots. In: Semisentient Robots, *IEEE Intelligent Systems*, 16 (5), pp. 16-22.

Steels, L., Vogt, P. 1997. Grounding adaptive language games in robotic agents. In: P.Husbands & I. Harvey (eds.), *Proc. ECAL97*, MIT Press/Bradford Books, Cambridge, MA, pp. 474-482.

von Uexküll, J. 1928. *Theoretische Biologie* (Theoretical Biology), Frankfurt/Main, Suhrkamp.

Wittgenstein, L. 1953. *Philosophische Untersuchungen. (Philosophical Investigations.)* Frankfurt am Main: Suhrkamp