From pre-representational cognition to language[1]

Takashi Ikegami and Jordan Zlatev

Abstract

Keywords: representation, pre-representational cognition, language, complex systems, active perception, dynamical categorization, sensory-motor coordination, meaning space, mimetics

1. Introduction

The issue of the nature and role of representations in cognitive science is a heavily contested one, with e.g. Johnson and Rohrer (*this volume*) arguing for a non-representational account of cognition, including language, while e.g. Gärdenfors (*this volume*) claims that his theory of "conceptual spaces" is fundamentally representational.

While "mental representation" was the most fundamental concept of "classical" cognitive science (e.g. Fodor 1981), the 1990s witnessed the rise of "second generation" cognitive science (e.g. Varela, Thompson and Rosch 1991) making heavy use of notions such as *embodiment* and *interaction* and this approach reacted against what was perceived to be an overextension of the term 'representation' to involve just about any kind of cognitive process. For example, Johnson and Lakoff (2002: 249-250) point out that: "As we said in *Philosophy in the Flesh*, the only workable theory of representations is one in which a representation is a flexible pattern of organism-environment interactions, and not some inner mental entity that somehow gets hooked up with parts of the external world by a strange relation called 'reference'."

However, there is a serious problem for such interactionist accounts to the extent that they purport to provide an explanation of language, since sensorimotor interaction is a non-representational notion (in any sensible

[1] *In press. Please don't quote without permission*.

use of the term 'representation'), while language *is* representational in two different, though related, respects: it has expression-content structure and statements are *about* (real or imagined) states of affairs (see Zlatev *this volume*; Sonesson *this volume*).

In this chapter, we hope to clarify this problem, and perhaps even offer some ingredients to its resolution. The first step involves conceptual analysis, in which we explain how we understand the concept of representation, and in particular, *mental* representation. We will distinguish between *pre-representational* and representational cognition, and point out their respective properties. We will suggest that only *true* representations are properly regarded as *concepts*.

In sections 3 and 4, respectively, we will characterize structures of pre-representational cognition: *dynamical category* and *internal meaning space*. A dynamical category is an emergent category resulting form sensorimotor interaction with the environment. Unlike explicit, "classical" concepts it can not be characterized in terms of necessary and sufficient conditions. In these two respects it resembles two of the central concepts within Cognitive Linguistics: *image schemas*, as defined by Johnson (1987) and Johnson and Rohrer (*this volume*) and *prototypes* (Rosch 1973; Lakoff 1987). However, the concept of dynamical category is most closely related to Gibson's (1979) ecological psychology (Costall *this volume*) and in particular to what is currently known as *active perception*.

By internal meaning space we mean a cross-modal network of dynamic categories. Similar to its constituent notion of dynamical category, the notion emphasizes the first-person, or subjective nature of cognition. Sensory input is meaningful for the organism due to its internal (neural) dynamics (Freeman 2003), which depend on an intrinsic value system (von Uexküll 1940; Zlatev 2003). At the same time, the notion of internal meaning space captures the fact that dynamical categories are not independent but linked in a "space" of different sensory modalities and dimensions. We will suggest that such a dynamic, internal, value-laden space serves as an essential intermediary between simple, reactive sensory-motor cognition and representational cognitive functions, including language. It resembles the "conceptual spaces" of Gärdenfors (2000, *this volume*), but differs in being dynamic and value-based, and thereby affect-laden. We will argue than an internal meaning space is implied in the phenomenon of *synesthesia* (Cytowic 2002).

In section 5, we suggest how such a pre-representational meaning space can help account for two aspects of language, *synesthetic metaphors* such as *sweet smile* and Japanese *mimetics* such as *suta-suta* ('walking hurriedly'). We follow the analysis of Kita (1997, 2001), who has argued that mimetics involve an "affecto-imagistic dimension" that is distinct from the "analytic dimension" that dominates linguistic meaning. Thus, the meanings of mimetics can be related in more than name to the concept of *mimetic schemas* (Zlatev 2005, *this volume*), which constitute pre-linguistic mental representations deriving from bodily imitation. In brief, while we will argue that language cannot be *reduced* to structures of pre-representational cognition, we will also show that such structures may be necessary for a complete account of it.

The concepts of dynamical category and meaning space may be rather difficult to grasp, especially if they are only characterized through language, which after all is not optimal for talking about pre-representational, dynamical phenomena with fuzzy boundaries. That is why our approach will be to use *complex systems* (Port & van Gelder 1998; Kaneko & Ikegami 1998; Kaneko & Tsuda 2000) in order to elucidate the nature of structures of pre-representational cognition. In particular, we will describe a number of computational simulations involving "artificial creatures" which "live" in a simulated "environment" and interact with it, whereby dynamical categories and a pre-conceptual meaning space emerge. To illustrate, we will show how even abstract phenomena such as triangles and rectangles can be categorized pre-representationally via blind touching, so that the criterion for their classification is defined by the creatures' styles of *action patterns*, rather than by their detached *concepts about* the figures.

We hasten to note, however, that these experiments should be regarded as a form of "Weak AI" rather than "Strong AI": there is a radical difference between *simulating* cognition and life and *duplicating* it (Searle 1992; Zimeke 2001). The "creatures" that are described in this chapter are not truly animate; they do not have any intrinsic value system, and thus have no basis for intrinsic intentionality and phenomenal experience (Zlatev 2003). We ask the reader to please remember this, since it will be tedious to have to point out again and again that when we refer to "creatures" in connection with the simulations described in Section 3, we are not succumbing to animism: the attribution of life properties to inanimate matter. Nevertheless, we believe that there are sufficient structural parallels between the mode of operation of models based on complex systems and real organisms to make the method justified as an analytical tool for studying "embodied" cognition.

2. What is a "mental representation"?

In a sense, the answer to the above question is rather straightforward: a representation is a structure that consists of *expression* and *content.* A clear example is a picture: the depicted apple cannot be eaten, but it represents (in this case *iconically*) an apple that can. The painting itself is the expression, and it is different from, at the same time as it corresponds to, something else. Whether this "something else" is a real specific apple, a generic apple, or an imagined apple is not important here: what is important is the *expression-content structure* itself. What "connects" the expression and the content is a process of *interpretation*: representations do not exist by themselves, but only for *someone* (cf. Sonesson *this volume*).

The real controversies begin when we ask whether there are *mental* representations, and if so what are they. "Classical", first-generation cognitive science found them everywhere: in thought, in language, in perception, in practical action (Gardner 1987). Some representatives of second generation, "embodied" cognitive science (Varela, Thompson & Rosch 1991; Johnson & Rohrer *this volume*) appear to take the opposite extreme and more or less abolish them.

We believe that these extremes are equally mistaken, and a golden mean is the best answer, in line with phenomenology (Husserl 1917) and Piaget's (1945) account of the origin of "symbols" in childhood: to say that a subject has a mental representation is to say that he can (a) *differentiate* between the expression and content, and (b) see the first as *corresponding* to the second, pretty much as in the picture example, mentioned above. Sonesson (1989, *this volume*) uses this as the major criterion to distinguish between (true) *signs* and simple *indexicalities* (based on contiguity and factorality, i.e. part-whole) and *iconicities* (based on resemblance), which are not signs but only a *ground* for indexical and iconic signs, respectively. One may (and usually does) ask: Who is the "someone" doing the differentiation and finding the correspondence? There are three types of answers, only the third of which will do:

(a) An unconscious processor or a "homunculus" in the head; that leads to infinite regress: we need to account for the ability of the homunculus to "see" the expression and content, and "figure out" that the first stands for the second, and then we need to account for the mental representations in *its* head, *ad infinitum* (Edelman 1992).

(b) The expression is a "symbol" that is associated with a "meaning" for someone *else* than the system that is actually using the symbols, as in a computer. Yes, but then the representation is not *intrinsic* to the system, but to the programmer, or whoever else is doing the "interpreting" (Searle 1992).

(c) The *subject* himself (or herself), i.e. the conscious individual who "owns" a (possibly internalized) mental representation. As when you close your eyes, and *imagine* an apple: you do not confuse your imagined apple with the one it represents, you differentiate between the two. On the other hand, when you see an apple, you do not think of your perception as a "representation" of an apple: you see the apple itself. So there are no mental representations in perception, but only in imagination (Piaget 1945; see also Zlatev *this volume*; Sonesson *this volume*).

Mental representations, as here defined, make an irreducible reference to *consciousness*. The notion of consciousness is, of course, vexed with even more riddles than that of representation, but since it has become again a (scientifically) respectable topic over the last 20 years,[2] considerable advances have been made in both philosophical discussions, e.g. of the "hard problem" of the irreducibility of qualitative experience (Chalmers 1997), in distinguishing between different kinds of consciousness, e.g. *affective* from *reflective* consciousness, as well as in understanding the neural underpinnings of this admittedly rather ellusive phenomenon (e.g. Edelman 1992; Damasio 2000). Mental representations involve both types of consciousness mentioned above: they are reflective, since they can be *accessed* and *thought about* independently of whatever they represent, but they are also affective since they have what phenomenologists call a "phenomenal contour" (Thompson 2001), a particular "flavor" (to use a metaphor), due to the intrinsic value system of the subject as a living being (Zlatev 2003). The "flavor" of my imagined apple is different from that of an imagined rotten banana.

Finally, since mental representations can be (in principle) accessed, unless they are "repressed" due to Freudian or other reasons (Searle 1992),

---

[2] As testified by a number of journals including *Journal of Consciousness Studies*, *Consciousness and Cognition,* and *PSYCHE: An Interdisciplinary Journal of Research on Consciousness* and annual conferences such as *Towards a Science of Consciousness (TSC).*

they can be themselves represented, or expressed – in language, pictures, gestures or some other external medium (more or less accuractly). That is why they can be said to constitute our *declarative* knowledge, as opposed to procedural skills such as bicycle-riding, which are not based on mental representations (contrary to the claims of first-generation cognitivists). The latter distinction is made clear by Mandler (2004), who consitently distinguishes between the two sorts of knowledge in her recent monograph, pointing out some of their respective characteristics:

> Procedural knowledge, both perceptual and motor, is inaccessible to consciousness. … In spite of taking in lots of information at once … it is also relatively slow to learn, and learning is accomplished by associative strengthening, typically over a number of trials, as in operand conditioning or perceptual schema formation. It aggregates frequency information. …
>
> Declarative or conceptual knowledge, in contrast, is accessible to awareness and is either describable in language, or, with a little analytic training, by drawing. It requires attention to be encoded into this format; this means that it is selective. … The system can learn information in a single trial (in small quantities, of course) simply by being told. In comparison to procedural knowledge, it is relatively context-free. (Mandler 2004: 55)

Similarly to Mandler (2004), we consider only declarative knowledge, which consists of mental representations, to be *conceptual*. Pre-representational cognition is no-less, and probably even more important for our survival, but it needs to be clearly distinguished from conceptual knowledge, and hence we will say that its major "building blocks" are sensorimotor schemas and categories, but not concepts.

We can now explain the prefix "pre" in our title: Pre-representational (pre-conceptual) knowledge preceeds representational cognition in (a) *phylogenesis*: all animals have it, but it is only certain that human beings have mental representations – though it is possible that at least some animals such as the great apes have some forms as well, e.g. dyadic mimesis (Zlatev *this volume*), (b) *ontogenesis*: the sensorimotor cognition of the young infant is pre-conceptual in the sense of Piaget, though not as long as Piaget assumed: at least 9-month old babies can *recall* past experiences, and that implies representations (Mandler 2004) and (c) *microgenesis*: pre-representational schemas and categories typically operate faster than conscious representational thought, and continue to serve as a constant basis for behavior, even when we have progressed past the sensorimotor period.

Table 1 summarizes some of the properties of representational cognition and mental representations, contrasting these with properties of prerepresentational cognition, along a number of dimensions.

Table 1. Comparison of mental representation and pre-representational cognition along a number of dimensions.

Dimension	*Mental representations*	*Pre-representational cognition*
Expression-content differentiation	+	-
Consciousness	Accessible	inaccessible
Major functions	thought, recall, planning	perception, recognition, self-motion
Learning	faster	slower
Evolution	later	ealier
Ontogeny	later	earlier
Guide to behavior	slower	faster
Type of memory	declarative	procedural
Status	conceptual	pre-conceptual
Example structures	mental images, mimetic schemas, symbols	interaction patterns, sensorimotor schemas

As shown in Table 1, the two sorts of cognition have *complementary* properties, and an account of human cognition, and possibly even that of certain “higher” animals, involves both types and their coordination. While we have here emphasized their differences in order to make a clear conceptual distinction, we need to mention two caveats, less we be accused of “dichotomizing”. First, since mental representations are evolutionarily and ontogenetically later than pre-representational cognition, they can be said to “emerge” from it, and hence will not be completely independent from their pre-representational, “bodily” roots. Second, and as implied above, any particular form of “higher cognition”, such as language use, will inevitably involve pre-representational, as well as representational structures and processes.

Neverthessess, we emphasize that mental representations cannot be reduced to pre-representational cognition: they constitute a qualitatively new ontological level, and in this sense, we are in disagreement with the claims of strong evolutionary continuity of e.g. Johnson and Rohrer (*this volume*). Language is based on mental representations, but involves yet one higher ontological level: that of consensual social reality, mutual knowledge (Popper 1962; Itkonen 1978, 2003; Zlatev this volume). At the same time, it is not completely independent from the pre-representational, sensorimotor roots of all cognition. In the remainder of this chapter, we will show some evidence for this by first explicating two related forms of pre-representational cognition, dynamical categories and internal meaning space, and then show how they "map" onto certain aspects of language.

3. Dynamical categories

The concept of dynamic categorization originates from Gibson's (1962, 1979) "ecological" theory of perception (see also Costall *this volume*), and in particular his emphasis on perception as a form of activity, which is currently often referred to as *active perception*. Gibson's insights have been developed and re-interpreted in multiple ways. Sasaki (2000, 2002) applied a Gibsonian analysis to various situations such as a blind man's navigation patterns in a town, people's usage of the visual landscape and the action structure of breaking an egg. According to Costall (*this volume*), the most important characteristics of Gibson's concept of *meaning* is that it is neither equivalent to external sensory input, nor to "representations" generated in one's brain, but constitutes a dynamical, relational category that arises as an active perceiver interacts with an environment.

The following aspects are particularly important in re-thinking the concepts of active perception and meaning for the present chapter (see Sasaki 2000, 2002; O'Regan & Noë 2001):

i) Perception can emerge via self-movement.
ii) Perceiving the environment means to explore it.
iii) Any action has inherent multiplicity.

As instance of the first aspect is *active touch*. Gibson (1962) reports on experiments with blind subjects touching different shapes of cookie cutters.

If the cutter was placed on the subjects' palms, they could only tell the correct shape with 50% accuracy. When the cutter was pushed randomly on one's palm, the subjects could tell with 72% accuracy. Only by touching the cutter in a self-guided manner could they recognize the object in more than 95% of the cases. This study also illustrates point (ii) that perception is a form of exploration: As we will discuss in this section, exploration is not just a method to arrive at perception; rather perception is *equivalent* to the on-going exploratory process.

The third aspect is especially important for simulations involving "artificial creatures". An issue that often comes up in this context is how to select the most appropriate set of actions. However, aspect (iii) implies that no discrete action set (a "plan") needs to be prepared in advance. Our body schema (see Gallagher 2005, *this volume*) has a huge number of degrees of freedom. Even a simple action pattern (e.g. sitting down on a chair) consists of multiple sub-level actions and the exact sequence will be afforded by the environment as the action proceeds.

The following example could perhaps clarify the distinction between "abstract" representational knowledge, and the kind of pre-representational know-how that is gained through active perception. In Japan it is customary for children to make paper cranes. We can learn how to make a paper crane by just looking at a picture of the end state of the process, with step-by-step instructions on how to get there (as in an IKEA manual). But this kind of knowledge is qualitatively different from that gained by actually making a paper crane. One's experience of making the paper crane, the way we fold the paper, how we feel touching origami and hear its rustling, etc.: all these complex perceptive experiences organize our "embodied" dynamical category of a paper crane. While folding origami, we experience trail-and-error everywhere giving rise to an exploratory process. Just by seeing and memorizing an algorithmic instruction we cannot form this category. So not only can representational cognition not be reduced to pre-representational cognition (as argued in Section 2), bodily skills cannot be reduced to (mental) representations either (Dreyfus & Dreyfus 1986).

3.1. Experiments with "artificial creatures"

The notion of active perception is often evoked within the "embodied cognition" approach in the field of *cognitive robotics* (Pfeifer & Scheier 2001). In this context "embodiment" refers to the spatial/temporal dynamics and the physical constraints of the robot's body with respect to a given environment. This is what Ziemke (2003) apparently means by "organismoid embodiment", which is similar to, but nevertheless different from the kind of embodiment that characterizes living systems ("organismic embodiment"). Since phenomenology presupposes biology, i.e. only (some) living beings are capable of qualia (Zlatev 2003; Emmeche *this volume*), there is no reason to suppose that such "artifical creatures" have any kind of phenomenal experience. Nevertheless, they can illustrate how categorization can result form mastering sensory-motor coordination with certain aspects of the environment, and in the process creating dynamical categories of these *without* forming concepts.

Recently, there have been many studies of dynamic categorization involving "artificial creatures", some of which are summarized in Table 2 below. Some of these studies involve actual, physical robots (Scheier & Pfeifer 1995), while others involve computational simulations (Morimoto & Ikegami 2004). While some scholars such as Brooks (1999) and Steels (1994) insist on the importance of using real physical devices in "embodied AI", there are good reasons to rely on simulations as well. As Ziemke (2003: XX) points out, "instead of focusing on one experiment or a few experiments with a real robot, many questions are more suitably addressed through a large number of experiments allowing for more variations on agent morphologies, control architectures, and or environments." Furthermore, it should be remembered that no matter the type of the experiment, such studies involve models which *simulate* rather the duplicate the phenomenon being studied (Searle 1992; Zlatev 2003), and this is more easily forgotten if the model is a "real" 3-dimensional, physical structure.

Another thing that is often forgotten, or at least not explicitly emphasized, but which is central for our discussion, is the difference between the (models of the) *categories* of the artificial creatures, and the corresponding (human) *concepts* such as LARGE/SMALL. The two are quite different: the first are "personal" and pre-representational, the second are (potentially) shared and representational. What the models simulate is categories of the first type, and it is only we as external observers possessing the respectve concepts who can see correspondences between the columns in Table 2.

Table 2. Studies of computer simulations of dynamical categorization. The first column shows the corresponding categories that the “creatures” formed on the basis of physical motion. The second column shows a specification of the differences in conceptual terms. It is only our interpretation that links the two.

Dynamical categories	*Concepts*	*References*
GRASPING/NON-GRASPING	LARGE/SMALL	Scheier & Pfeifer (1995)
NAVIGATION DIFFERENCES	ROOM A/ROOM B	Tani & Nolfi (1998)
NAVIGATION DIFFERENCES	TRIANGLE/RECTANGLE	Morimoto & Ikegami (2004) Marocco & Floreano (2002)
FINGER MOTION DIFFERENCES	SPHERE/CUBIC	Nolfi & Marocco (2002)
APPROACHING/AVOIDING	FAST/SLOW BLINKING	Iizuka & Ikegami (2004)

To understand this better, we offer a brief summary of how the “creatures” in the first two studies of Table 1 performed classification through *dynamical categorization*. It is worth having in mind the similarities with Gibson’s (1962) psychological experiment involving blind touch mentioned earlier: in both cases perception is performed through self-guided motion and exploration.

Pfeifer and Scheier (1995) conducted studies of robots that learned to discriminate on the basis of the size of an object by their bodily movements. In the experiment, there were pegs of two different sizes, one with a large diameter and one with a smaller one, distributed in an arena. A robot used a gripping wire to pick up pegs after sensing these with its light sensors. The robot could pick up the smaller pegs to carry them over to its nest but not the larger ones. Using (a model of) an unconscious learning process, it came to discriminate small and large pegs. As learning proceeded, the robot neglected large pegs and only tried to pick up the smaller ones. On the basis of this, one could be tempted to say that the robot came to *know* the concepts of LARGE and SMALL (relative to its own “embodiment”), but that would be a mistake. Rather, by mastering a specific type of sensory-motor coordination, the robots learned the categories GRASPABLE/NON-GRASPABLE, as they apply in this particular context. For example, if we changed the surface texture of the pegs so that some “small” ones would be smooth, they would become difficult to grasp, and would therefore be categorized together with the large ones, showing again the difference between categories and concepts.

Another experiment demonstrated that robots could discriminate between two different rooms by their bodily movements (Tani & Nolfi 1998). The two rooms were distinguished due to their different spatial arrangements of corners and walls. A door that connected the two rooms randomly opened and closed. When the door opened, a robot could move into the other room by chance. After going back and forth between the two rooms, a robot came to "know" which room it is situated in: it came to possess two mutually exclusive neural modules, one for room A and one for room B, and depending on which one was activated it behaved differently. We can say the robot categorized some of the spatial characteristics of its environment, but can we say that it had *concepts* of the two rooms? The distinctions made in Section 2 lead to a negative answer. The "neural modules" corresponding to the two rooms are not mental representations of the two, since there was no way for the robot to differentiate between the module and (its perception of) the room. There was therefore no way for it to *imagine* room B, while it was in room A, and to "decide" to go there.[3]

On the other hand, having concepts of the two rooms implies being able to link the concept to specific perceptual details such as the color of the walls or the different light pattern around a corner. Concepts must be "grounded" in perception – or else they are "empty", as Kant famously pointed out. Our proposal is that dynamical categories play an important role in this grounding.

3.2. Dynamic categorization of object shapes: a case study

In order to gain a better understanding of how such models of dynamical categorization work, and to appreciate both the strengths and the weaknesses of dynamical categories, we here describe in some more detail the study of Morimoto and Ikegami (2004), in which simulated artificial creatures could learn to classify triangles and rectangles through self-guided motion. The goal of the experiment was to see if geometrically well-defined concepts such as TRIANGLE and RECTANGLE could be classified only by local "blind touch" and "exploration".

[3] The classification of the two rooms was also dependent upon such contingencies as the dynamics of the door between the two rooms. Indeed it turned out that the open/shut dynamics of the door that connected room A and room B was essential: randomness was needed to achieve successful classification (Tani p.c.).

The “creatures” of Morimoto and Ikegami (2004) “live” in a 2-dimensional world, populated by objects of various shapes and colors, as shown in Figure 1. The objects are 2 different kinds of triangles and 4 different kinds of rectangles. As a creature “explores” the environment, it “touches” some of these objects, and then changes its style of motion. Since the creature’s receptive field is highly limited, it can not perceive the shape of a whole figure. The creatures were “evolved” using a genetic algorithm, which mimics biological evolution (Nolfi & Floreano 2001), and after thousands of evolutionary generations, using a “fitness function” that favors creatures that avoid triangles and explore rectangles, some creatures succeeded in spending more time in touching rectangles and less time in triangles. The point is that the creatures were not in any way explicitly trained to distinguish triangles from rectangles and the classification was established only as a result of the actual exploration.

Figure 1 shows part of the environment with (a) several triangles and rectangles randomly distributed and (b) narrow lines penetrating the objects. The lines represent a creature’s *spatial trails* and we can see that it by-passes triangles but “explores” rectangles, the black part representing points contact between creature and object. In particular, it can be seen that the creature has densely covered a rotated square in the top-right corner.

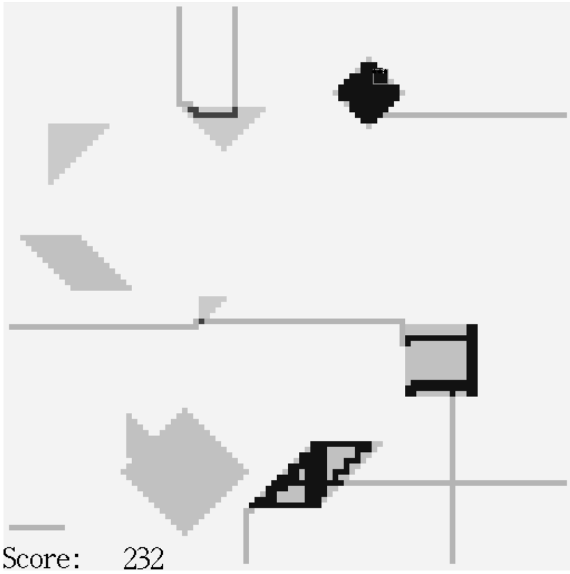

Figure 1. Part of a 2-dimensional environment showing (a) rectangles and triangles and (b) the spatial trail of a "creature", shown as a line. The black parts of the figures represent points of touch between creature and objects (for details, see Morimoto and Ikegami 2004)

How did the creatures achieve this implicit classification? Each creature was controlled by a simple artificial neural network, shown in Figure 2, of the type used in many other similar models (e.g. Ziemke and Thieme 2002). The network receives sensory input in a 3x3 matrix of "neurons", each one of which has a continuous value from 0 to 1.

A set of connections projects from the sensory input to the motor output, both directly, and via a number of "context neurons". Each input neuron has "synaptic" connections to the other input neurons, and similarly for the context neurons. Each neuron receives an integrated signal from the other neurons multiplied by the connection weight. The connection weights also take continuous values, and are initially set by the genetic algorithm. The integrated signal filtered by a sigmoid function gives a new neural state of that neuron. This updating schedule (neural states at time T > integration of weighted activities > neural states at time T +1) is recursively iterated while a creature is moving around the environment. The output signals from the input neurons are integrated to produce motor outputs. There are three output neurons (L, F and R). The most active output neuron determines the next action of the creature: turn right (R), go straight (F), or turn left (L). The weights of the connections were not fixed, but adapted by a process of Hebbian learning (Hebb 1949), to allow predicting the sensory stimuli of one step in the future. This plasticity enables dynamic switching from one navigation mode to another: "wandering", "exploring objects" and "filling in" (for details, see Morimoto & Ikegami 2004).

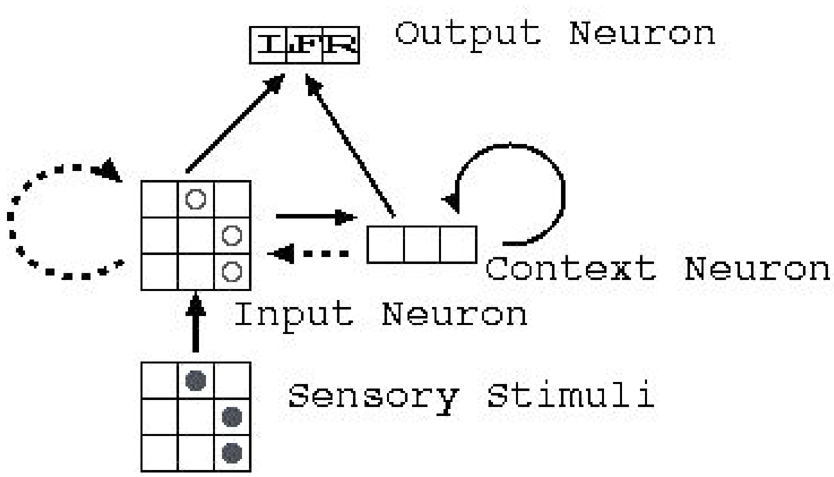

Figure 2. The design of the neural network that controlled the artificial creature's "behavior" (Adapted from Morimoto & Ikegami 2004)

Two important observations concerning the classification and the way in which it was achieved need to be made. First, the creatures could not learn to discriminate between triangles and rectangles completely accurately, but rather formed *prototype-like categories*. Figure 3 shows an example of a creature's "behavior" towards both different instances of the same object and different objects, given the same initial neural state (i.e. connection weights and neuron activations). It can be noticed that a creature's behavior looks very different depending on where the creature first encounters the object:

a) The first 10 instances of Figure 3, starting from the top, are triangles. As can be seen, the creature enters and leaves these objects without spending much time exploring them.
b) The following 4 instances are squares, and the creature correctly

"fills in" these figures more or less completely.

c) The following 6 instances are all parallelograms. The creature fills in the insides of the first three quite thoroughly, but not the other three.
d) The remaining 12 instances are trapezoids. The creature's behavior depends on the orientation of the trapezoid: the first 6 instances are explored but their reversed images are not, and are thus not differentiated from triangles.

Figure 3. A creature's "imperfect" categorization of triangles (the first 10 fugures) and rectangles (the remaining figures). (See the text for discussion)

Thus, we can conclude that the creature treated squares as the most prototypical type of rectangle, while other rectangular shapes were not consistently distinguished from triangles.

The second observation is that the classification is highly *context-sensitive*. Different instances of the same object were treated differently depending on (a) the creature's internal state and (b) the spatial arrangement of the objects. For example, the order in which the objects were explored (rectangle to triangle or triangle to triangle) changes the internal state of a creature so that the same triangle is "perceived" differently (at different points in time).

The key to understanding these features of the model, as well as its relative success in the discrimination task, were the *internal context neurons*, which functioned as a sort of procedural memory. As shown in Figure 2, the creature's behavior was not simply a product of the coordination between bare sensory input and motor output, but also involved *processed* sensory input that was "stored" in the context neurons and the weights of the connections to and from these units. Specific analyses showed that a creature could be sometimes driven by the sensory input (the left circuit) but sometimes by the context neural states (the right circuit), see Figure 2. The conclusion is that in order to perform (prototypical, context-sensitive) triangle/rectangle discrimination, the creature *required the combination of "raw" input signals and the internal neural states.* Without this combination, the creatures showed much less diverse behavior and lower performance in discrimination.

3.3. Further studies of dynamical categorization

In order to further explore the nature of dynamical categorization and specifically the role of *internal* dynamics, Iizuka and Ikegami (2004) conducted experiments with a new neural network architecture, where "creatures" were able to spontaneously turn on and turn off the signal from the sensory input. If the signal was "off" then behavior was guided soley by the internal dynamics (i.e. memory) of the system, which also controlled "autonomously" whether environmental input should be taken in or not. This can be compared to a mechanism that differentially focuses attention on either the environment or on the memory of past experience. Since such differentiation is the first step towards representation, as argued in Section 2, the model of Iizuka and Ikegami (2004) could be seen as a model of a *step* towards mental representation. But only a relatively small one, since there is of course no awareness of any correspondence between the internal state and the environment.

The model was applied to a different task: differentiating between different frequencies of light blinking by approaching, respectively avoding the light source. It turned out that the model learned the appropriate behaviors much more easily with the help of the "selective attention" mechanism than without it. The reason was that by being able to selectively direct its attention either to the blinking patterns or to its memory of past interactions, a creature's behavior was not determined so much by the external stimulus as by the style of its own motion in previous similar circumstances.

In sum, the experiments with "artificial creatures" described in this section show how dynamical categories can emerge as result of sensorimotor interactions with various stimuli, without there being any mental representations (i.e. concepts), or even awareness of these stimuli. We showed that such categories display features such as prototype effects and context-sensitivity, which may be advantageous in some circumstances, but disadvantageous in others. In particular, we suggested that dynamical categorization can be enhanced through internal "context units" or other mechanisms implementing a sort of memory, and even more so by the ability to selectively pay attention to the internal states rather than only to the external environment. We use this as departure point for a discussion of the role of these internal states, or what we call internal meaning space.

4. Internal meaning space

The experiments described in the previous section illustrated some Gibsonian features such as the active nature of perception and the "multiplicity of action", and how they could emerge from a combination of evolutionary history and a history of spatial exploration. At the same time, we pointed out that the dynamic internal state of the simulated creatures retains a form of long-term memory or personal history (which is something underplayed in Gibsonian theory), and this affects the selection of the current action pattern. This would imply that a living creature with such a memory would experience the world differently from creatures with different *histories* of structural coupling (Varela, Thompson and Rosch 1991). In this sense, we can say that a creature with an internal state is not simply coordinating sensory input and motor output, but coordinates the *meaning* of the sensory input with self-motion: it "interprets" its environment, rather than just reacting to it.[4]

In section 2 we argued that there is a gap between pre-representational cognition and (true) mental representations. The experiments with "artificial creatures" helped clarify this gap, but also showed that the *internal dynamics* of the control architectures that govern the behaviors of the creatures were essential for the formation of adequate dynamical categories. We will here argue that such internal dynamics constitute an *internal meaning space.* This concept does not "bridge the gap" between sensorimotor and representational cognition, but nevertheless constitutes an important (evolutionary) step in the direction of mental representations.

4.1. Quality dimensions, comparison with Gärdenfors (2000)

Let us explain the senses in which we are using the constituent terms of the notion *internal meaning space*. The fact that it is "internal" does not mean that it is representational, but only that it is *subjective*, an emergent categorization of the environment performed by the organism, its Umwelt (von Uexküll 1940, see Lindblom & Ziemke *this volume*). The "meaning"

---

[4] On the basis of his long-term animal experimental studies, Freeman (2003) insists that neural activity patterns in the sensory cortex reflect the meaning of the input rather than the actual stimulus. See also the discussion by Ziemke & Sharkey (2001) concerning von Uexküll's notion of the "historical basis of reaction".

of the space lies in the fact that it consists of a network of dynamical categories which are value-laden. But in what way is it useful to regard it as a "space"?

This metaphorical expression becomes justified if, similarly to Gärdenfors (2000, *this volume*), we regard the internal meaning space to be structured by a number of *quality dimensions*, such as HEIGHT, BRIGHTNESS, PITCH, VOLUME, WEIGHT and SHARPNESS. These are not just abstract dimensions along which stimuli can be classified, but dimensions that *matter* to living creatures, in the sense of supporting their self-preservation and are hence intrinsically meaningful (Zlatev 2003). Also similar to Gärdenfors, we would claim that some of these dimensions have been selected for in evolution and are thus in one sense of the word "innate", while others are learned, or at least modulated, as a result of experience. A final similarity is that we find Gärdenfors's notion of "natural concepts" as *convex regions in conceptual space* useful and believe that it corresponds to our notion of dynamical categories: both allow for prototype effects, with the prototype corresponding to a particular point, or as we would prefer to say, to an *attractor* in the internal meaning space.

The differences between our concepts and those of Gärdenfors are basically two. First, the notion of internal meaning space, with dynamic categories as attractor states, is considerably more dynamic than what Gärdenfors' model allows (even with the extensions of his model involving forces): dynamical categories are *constructed* through sensorimotor activity in the manner illustrated by the simulations in Section 3, and do not simply exist as static "convex regions". Second, we maintain that Gärdenfors conflates pre-representational internal meaning space and true representations by calling both "representations". This is shown clearly in the quotation that he refers to as well:

> Evidence suggests that dimensions that are easily separated by adults such as brightness and size of a square are treated as fused together by children. … For example, children have difficulty identifying whether two object differ on their brightness or size, even though they can easily see that they differ in some way. Both differentiation and dimensionalization occur throughout one's lifetime. (Goldstone & Barsalou 1998: 252)

We interpret this as evidence that the dimensions of 2 year-old children are still pre-representational, and that only when they develop the capacity to differentiate between the dimensions *consciously*, or what Gärdenfors (*this volume*) calls "to reason about the dimensions", do they become truly representational, or conceptual. Similarly, experientially relevant "regions" in internal meaning space are for us natural *categories*, but not (yet) concepts. This may sound like terminological hair-splitting, but it is not, since we maintain a difference that Gärdenfors does not. This leads Gärdenfors (*this volume*) into difficulties when it comes to distinguishing between forces as "psychological constructs" and as "scientific dimensions", and to a serious misrepresentation of the *mimetic schemas* (Zlatev 2005, *this volume*) as "patterns of forces… represented in the brain". Most crucially for the purpose of the chapter, we emphasize the pre-representational nature of the internal meaning space since we are interested in clarifying, and to some extent bridging the gap between non-representational cognition and language, and we suggest that the internal meaning space plays an important role in this respect.

4.2. Cross-modal transfer and synaesthesia

An important property of the internal meaning space is that it is intrinsically *cross-modal*. Indeed, Edelman (1992) argues that at least two modalities ("a classification couple") are necessary for any kind of natural categories to emerge. While the experiments described in Section 3 were multi-modal in the sense that they involved simulations of coordination between haptic sensation and self-motion, perceptual experience is much richer than that. The quality dimensions mentioned earlier derive from different modalities: vision (e.g. BRIGHTNESS), auditory system (e.g. PITCH), kinesthetic sense (e.g. WEIGHT) and haptic sense (e.g. SHARPNESS), and the major function of the internal meaning space is to provide an *integration* of the different modalities.

A good deal of this integration is due to experience, by coordinating the modalities through exploration, e.g. as proposed by Piaget (1945). But this cannot explain all the available data, since it appears that some forms of coordination between modalities, often described in terms of "transfer" or "mappings" are pre-established by evolutionary processes (analogous to the "genetic algorithm" in the simulation described in Section 3). Meltzoff and Borton (1979) showed that infants as young as 1 month looked longer at pacifiers that they had previously explored only orally, displaying a transfer between haptic sense and vision. Perhaps even more famously (and at first controversially), a number of studies (Meltzoff & Moore 1977, 1995) have shown that neonates are able to copy a number of bodily actions such as tounge protrusion and mouth opening, practically from birth (see Gallagher *this volume*). More recently, such displays of “blind imitation” (the infant cannot see its own body part) have been shown for chimpanzees as well (Myowa-Yamakoshi et al. 2004). These results imply mappings between visual perception of body motion and kinesthetic perception of one’s own body and have been linked to specific neural mechanisms such as “mirror neurons” (e.g. Rizzolatti & Arbib 1998). While not offering evidence of innateness, the perception and classification of types of *body motion* (Johansson 1973, see Gärdenfors *this volume*) through visually minimal information can also be explained as the result of mapping from vision to the *body-schema* (Gallagher 2005, *this volume*).

All these mappings, we suggest, take place in the internal meaning space. They do not require to be thought about, or inferred, but are “perceived directly” in phenomenological terms. This is in line with our claims that internal meaning space is pre-representational, in contrast to Gärdenfors (*this volume*), as well as Meltzoff and Borton (1979) who interpreted their results as showing that infants “represent” objects in an *amodal* representational format. Rather, we view the internal meaning space as *cross-modal*, and initially inaccessible to reflective consciousness.

A phenomenon that can be interpreted in terms of a cross-modal internal meaning space is *synaesthesia* (Cytowic 1995, 2002; Baron-Cohen 1996; Ramachandran and Hubbard 2001). Despite controvercies on the mechanisms responsible for it, there is now a consensus that “synaesthesia (Greek, *syn* = together + *aesthesis* = perception is the involuntary physical experience of a cross-modal association” (Cytowic 1995). Clinically, synesthesia is present in (at least) 1 in 20,000 individuals, with a higher rate for women than men (6:1) and is genetically inherited. Synaesthesia has a number of characteristics:

(a) it is involuntary and “unsuppressable”: the subject cannot help but, for example, see a certain color on hearing a particular tone,

(b) it can involve all modalities, but some are more common than others (see Section 5 below),

(c) it is is usually unidirectional, e.g. different sounds evoke visions, but vision does not typically evoke sound,

(d) it is “projected”, perceived externally, as ordinary perception, rather than in “the mind’s eye” and

(e) it is emotional: either disturbing or rewarding, but never neutral.

The involuntary character of synaesthesia is sufficient to distinguish it from mental representations, and in particular language. While language may involve some similarities with synaesthesia, as we will point out below, we emphasize again that language, as a conventional symbolic system for communication and thought involves (intersubjective) representations, and is qualitatively different from synaesthesia. To quote Cytowic (1995) again: "Its phenomenology clearly distinguishes it from metaphor, literary tropes, sound symbolism, and deliberate artistic contrivances that sometimes employ the term 'synaesthesia' to describe their multi-sensory joinings." On the other hand, synaesthsia is obviously not "out there" in the objective environment. In our terms, it is clearly pre-representational.

While we followed Cytowic in emphasizig its "distinctness" as a phenomenon, synaesthesia is possibly more continuous to ordinary perception than so, and in a sense "we are all synaesthetics" to some degree. Maurer (1993) argues, on the basis of both animal and human evidence, that neonatal cross-modal transfer is essentially synaestheic, but that it normally disappears after the few months:

> During early infancy – and only during early infancy – […] evoked responses to spoken language (are recorded) not just over the temporal cortex, where one would expect to find them, but over the occipital cortex as well. There are similar reports of wide-spread cortical responses to visual stimuli during the first 2 months of life. […] Results such as these suggest that primary sensory cortex is not so specialized in the young infant as in the adult. (ibid: 111, quoted in Baron-Cohen 1996).

If so, that could possibly explain the phenomenon of neonatal imitation, which we argued above is also a form of cross-modal transfer – as well as its "disappearance" after the first 3 months or life. However, cross-modal transfer "reappears" later in life, e.g. in the form of *mental simulation* (see Svensson and Ziemke *this volume*), but we would argue that these effects reflect imagination, rather than perception, and are thus representational. There is thus room for both discontinuity and *some* continuity between internal meaning space and representational cognition, including language, as we will suggest in the next section.

5. Representation, quasi-synaesthesia and language

What kind of change in internal meaning space is necessary in order to give rise to mental representations? This is a difficult question to which we will here offer a simple and preliminary answer: Representation involves a *bifurcation* of the internal meaning space into (a) perceptual consciousness and (b) imagination (reflective consciousness). The model of Iizuka and Ikegami (2004), reviewed at the end of Section 3, can serve as an illustration. To the extent that the subject can attend differentially to aspects of the space itself, separate from the way they mediate the perception of the external world, that subject will be capable of "mental images". What remains is to understand these images as actually *standing for* something else than themselves. With this the conditions for mental representations defined in Section 2 will be fulfilled.

How does this "bifurcation" come about? Piaget (1945) argued that *imitation* plays a central role for this in childhood. Zlatev (*this volume*) explains this in terms of *bodily mimesis*: initially the child does not differentiate between the perceived body motion of the other, and the motion of its own body, as in synesthesia. But with deferred imitation comes differentiation. And with "representational imitation" comes the ability to (consciously) access the mimetic schema, and use it as a model to guide future action. However, an apparent problem for these accounts is posed by children paralyzed, or severely motorically impaired, from birth (e.g. Jordan 1972). Hence, *actual* imitation cannot be a necessary condition for the emergence of mental representations. Zlatev (2005) attempts to deal with this in terms of *covert* imitation, but this does not answer where the ability to "covertly imitate" derives from.

At present there does not seem to be any good explanation, but we can here at least offer a description: to have a mental representation is to have one part of internal meaning space, imagination, standing in correspondence with another, perception, and being able to differentially focus attention to one or the other. This description, however, reminds of the phenomenon of *synaesthesia* reviewed at the end of the last section. The important differences are that (a) attention is under voluntary control and (b) the "synaesthetic experience" is not projected into the perceptual world, but *understood* as internal and "unreal" (in normal conditions).

If this conjecture is correct, we would expect to find "quasi-synaesthetic" experiences: imagined projections between modalities in the arts, as well as in language. In this section, we will briefly describe two linguistic phenomena which seem to provide some support for the view that at least some forms of linguistic representation may involve what we here call *quasi-synesthesia* in order to distinguish it from true synaesthesia as defined in Section 4.2.

5.1. Quasi-synaesthetic metaphors

What we will refer to as quasi-synaesthetic metaphors are more commonly known as "synaesthetic metaphors" (e.g Ullman 1964): expressions such as *sweet smile*, *cold look*, *soft music*, and *loud color*, where experience from one sense ("source domain") is projected to, or mapped onto a phenomenon that is primarily perceived in terms of another sense. For the expressions give above, that would involve: Taste → Vision, Temperature → Vision, Touch → Hearing and Hearing → Vision, respectively. While some of these expressions may be quite conventional, others may be more "novel", e.g. *bitter chuckles* (Tomas Pynchon, *Gravity's Rainbow*). The specific types of expressions are language- and culture-specific, but the phenomenon is apparently universal (Osgood 1959). Furthermore, many have observed that there appear to be some universal tendencies. Ullman (1964: 86) writes:

> [T]he movement of synaesthetic metaphors is not haphazard but conforms to a basic pattern. I have collected data for the sources and destinations of such images in a dozen nineteenth-century poets, French, English and American, and found three tendencies which stood out very clearly: (1) transfers from the lower to the more differentiated senses were more frequent that those that map in the opposite direction: over 80 percept of a total of 2000 examples showed this "upward trend"; (2) touch was in each case the largest single source, and (3) sound the largest recipient.

If these tendencies are found across languages and cultures, then that would appear to imply some general cognitive motivation of semantic structure, of the type often evoked in Cognitive Linguistics. At the same time, "mappings" from *less* differentiated domains (Taste and Touch) to *more* differentiated domain (Vision and Hearing) would appear to go against the grain of Conceptual Metaphor Theory (Lakoff & Johnson 1980, 1999), where rather the latter is expected: "The greater inferential complexity of the sensory and motor domains gives the metaphors an asymmetric character, with inferences flowing in one direction only." (Lakoff & Johnon 1999: 57-58)

Ullman does not specify his evidence in more detail, but Day (1996) reports a study in which he analyzed 1269 (quasi-)synaesthetic metaphors in English texts ranging from Chaucer to Pynchon, specifying the directionality of the mappings. Table 3, adapted from Day (1996, Table 7), classifies the 1269 expressions in terms of these mappings between the "six senses". It can be seen, for example, that Touch was by far the most common source domain, mapping to Vision 135 times (e.g. *hard look*), while Hearing was, also by far, the most common Target domain, exactly as claimed by Ullman (1964).

Table 3. Classification of the mappings in 1269 quasi-synaesthetic metaphors, based on Day (1996: Table 7), see text for discussion.

Target domain	**Source domain**						
	Hearing	**Taste**	**Smell**	**Temp.**	**Touch**	**Vision**	# *Target domain*
Hearing	-	149	1	86	540	80	856
Taste	0	-	0	1	6	0	7
Smell	7	60	-	3	34	14	118
Temp.	0	19	0	-	8	4	31
Touch	3	10	0	0	-	2	15
Vision	26	38	1	42	135	-	242
# Source domain	36	276	2	132	723	100	1269

When Day (1996) subtracted the instanses when a sense modality was used as source from those when it was used as target, he arrived at the “sensory ranking” given in (1):

(1) Touch > Taste > Temperature > Smell > Vision > Hearing

Analogously, Day classified the types of synaesthesia from the 25 subjectes reported by Cytowic (2002 [1989]), some of who had multiple synaesthesia, e.g. MMo linked both sound impressions to sensations of taste and vision, and visual impressions to sensations of touch, giving a total of 35 projections. These are given in Table 4, adapted from Day (1996, Table 6). Comparing to Table 3 shows that similarly, Hearing was by far the most common target (i.e. when the subjects heard sounds they also experienced sensations from other modalities, but unlike with the metaphors, it was Vision that was the most typical source, which is often stated in descriptions of synaesthsia, typically involving projection of *color* (e.g. Ramachandran and Hubbard 2001).

Table 4. Classification of 35 projections in synaesthetic 25 subjects, analyzed by Cytowic (1989), adapted from (1996: Table 5), see text for discussion.

Target modality	**Source modality**						
	Hearing	**Taste**	**Smell**	**Temp.**	**Touch**	**Vision**	*# Target modality*
Hearing		2	1	0	2	21	26
Taste	0		0	1	1	1	3
Smell	0	0		1	1	0	2
Temp.	0	0	0		0	0	0
Touch	0	0	0	0		2	2
Vision	0	0	1	0	1		2
# Source modality	0	2	2	2	5	24	35

Using the same method of “subtraction” provided the ranking given in (2), which, however, given the small number of subjects and projections shoud not be taken as too seriously. Nevertheless, it clearly shows that there are both similarities between quasi-synaesthetic metaphors and actual synaesthesia, above all that Hearing is most often “interpreted” as something else, and differences: above all the ranking of Vision.

(2) Vision > Touch > Temperature > Smell > Taste > Hearing

The conclusions that we draw from this are in part similar to those drawn by Day (1996): that while the similarities may reflect universal features of human consciousness, and its neural underpinnings, or in our terms, the organization of our internal meaning space, the differences require distinguishing between the mechanisms of synaesthesia and semantic processes: “The meanings for synaestheic metaphors are not simply there, hard-wired and innate, but are generated through semantic processes and fashioned by time and cultural elements, much like other metaphors” (ibid: 20). We further add that the differences point to the radical difference between pre-representational cognition and language. The first could possibly motivate some of its characteristics (e.g. the tendency to *consciously* relate sound-based concepts to concepts based on some other modality), but semantics cannot be reduced to pre-representational cognition, and even less so to neural structures.

Still, we find this line of research intriguing, and pointing to the possible synaesthetic roots of conscious, mental representations. Once the "source" and the "target" modalities can be differentiated, the first can be mapped onto the other, utilizing the quality dimensions of the internal meaning space. If the mapping between the perception of another's actions and those of oneself involves similar mechanisms to those involved in syneasthsia, its differentiation and focus on the "internal image" of the action would be equivalent to (dyadic) covert mimesis, possibly the original form of mental representation, as suggested by Piaget (1945) and Zlatev (2005, *this volume*).

5.2. Japanese mimetics

Sound-symbolism is another universal phenomenon of language, but the degree to which particular languages employ it varies. It is comparatively marginalized in Indo-European languages, and unsurprisingly, it was not considered a central feature of language from the birth of modern linguistics in Europe (Saussure 1916). Another reason for its de-appreciation is that it goes against the Saussurian dictum of the "arbitrariness of the linguistic sign". However, while language is *conventional*, this need not imply arbitrariness, even though these two concepts are often conflated (see Zlatev *this volume*). Japanese mimetics are highly conventional, but "aspects of the form meaning relationship are not arbitrary but are motivated by iconcity." (Kita 2001: 419-420). They are also a central feature of the language. Ivanova (2002: 2) provides the following informative charateriztion of their role in the language:

> Japanese is one of the languages with vast sound-symbolic systems […] with more than 2,000 onomatopoeic and mimetic words. These words overwhelm ordinary speech, literature and the media due to their expressiveness and load of information. Although they are never used in official documents, it is not exceptional to hear them in formal situations, too. People of all ages employ mimetic words in communication, believing that speech that abounds in such words sounds much more natural and full of life than speech that tends to avoid them.

Classifications of Japanese mimetics differ, but at least four types can be distinguished (Martin 1975; Kita 1997; Ivanova 2002; Baba 2003):

(a) *gi-sei-go*, words imitative of sounds produced by living creatures, e.g *wan-wan* ('bow-bow')
(b) *gi-on-go*, words imitative of sounds produced by the inanimate world, e.g. *ban* ('bang')
(c) *gi-tai-go*, words imitative of physical actions, e.g. *koro-koro* ('light object rolling repreatedly')
(d) *gi-jyoo-go*, words imitative of phychological states, e.g. *muka-tsuku* ('irritating')

The first thing that needs to be explained is the sense in which these words can be said to be "imitative", i.e. their iconicity. In the case of (a) and (b), which are similar to the onomatopoietic words that we are all familiar with, this is relatively straightforward: the sound-expression resembles the sound that is produced in the referential scene. Hamano (1998), who is often refered to as providing the most extensive analysis of Japanese mimetics, attempts to link particular phonemes (e.g. /p/ vs. /d/) and distinctive features (e.g. +/- voice) to specific meaning components. Thus, /p/ is for exampled claimed to be associated with "light, small, fine", /b/ with "heavy, large, course" and /m/ with "murkiness". However there are problems for this analysis. As Ivanova (2002) points out mimetic "words with initial /m/ are *maza-maza* (clearly, vividly), *meki-meki* (remarkably, fast), *miQchiri* (hard, severe), *moya-moya* (hazy, murky), *muka-muka* (retch, go mad). It is clear that "murkiness" is not their common semantic feature." Instead, she characterizes the expression-meaning correspondence of 199 mimetics of the (c) and (d) types in terms of more general "phonaesthematic patterns", where "phonaesthematic describes the presense of sequence of phonemes shared by words with some perceived common element in meaning" (Ivanova 2002). An example of such a pattern is given in (3) below.

(3) expression: *g/k* + V + chi + g/k + V + chi
meaning: LACK OF FLUIDITY OR SPACE
examples: *gichi-gichi* ('very tight'), *kachi-kachi* ('frozed hard', 'dried up'), *kochi-kochi* ('tense', 'stiff', 'flozed hard')

What this approach, however, leaves out are generalizations that apply within and accross the patterns, such as that contrast +/- voice may be used to distinguish WEIGHT, e.g. *koro-koro* vs. *goro-goro* 'heavy object rolling repeatedly' or VOLUME, e.g. *chara-chara* 'few coins rattling' vs. *jara-jara* 'many coins rattling' – as when the coins "come out of a slot machine when one hits the jackpot" (Baba 2003: 1868). Finally, neither of these accounts explains in which way the expression of the mimetic "imitates" or "resembles" their meanings, a problem that is even more pronounced with respect to *gi-jyoo-go*, sometimes also called "psychomimes". How can we make sense of the suggestion that their expression is mimetic with respect to psychological states?

We believe that a key to this puzzle is offered by the presence of cross-modal mappings in internal meaning space. In their investigations of synaesthesia Ramachandran and Hubbard (2001) describe experiments in which ordinary subjects were given contrasting pictures of objects of different shapes, for example one which was roundish and "soft" and another which was edgy and "sharp". Then they were given two "names", e.g. *kiki* and *bouba*, and asked to pair name and object. Just as expected, 95% of the subjects paired *kiki* with the sharp object and *bouba* with the roundish one.[5] Why should this be the case? If we start from the shapes, the cross-modal mapping between vision and touch would allow them to be *perceived* as "soft" vs. "sharp", motivating the use of these quasi-synaesthetic metaphors as a natural way to describe these figures. From the side of the expressions the production of the velar stop /k/, even more so combined with the front, unrounded vowel /i/ involves obstructions and narrowings in the vocal tract, which can similarly be perceived as "sharp" and "edgy". On the other hand, the shape of the vocal tract and the lips in the production of /u/ in *bouba*, are quite literally "roundish" and the passage of the air is "soft". The mappings between the senses Vision-Touch-Proprioception-Sound in internal meaning space thus provides for a correspondence between the shapes and the labels that would be impossible otherwise. A robot or a Martian with a very different kind of

[5] The experiment was essentially a replication of a classical experiment performed by Köhler (1929), who called the figures *takete* and *baluma*.

body (and possibly even a person lacking haptic sense and proprioception) would not be able to perceive the iconicity involved.[6]

Returning to Japanese mimetics, we can suggest that they are quasi-synaesthetic in a similar way to *kiki* and *bouba*. This can explain some of the contrasts, such as +/- voice, since voicing involves a higher degree of energy, both in terms of perception and production. But notice that even for *kiki* and *bouba* it is not possible to figure out what they would mean only on the basis of the iconicity of the expressions, but only to do the matching to the "correct" shape when one is provided a few shapes to choose from. Sonesson (2001) makes an important distinction between *primary iconicity*, in which the similarity between A and B can be perceived even without knowing that A is a representation of B, and *secondary iconicity*, in which to perceive any similarity between A and B requires knowing that A is a representation (sign) of B. Realistic pictures, photographs and pantomime can be interpreted by virtue of primary iconicity. On the other hand, the iconicity of diagrams in which "up stands for more" or onomatopoietic words like "bow-bow" can be appreciated first when one understands their representational expression-content structure. We would like to propose that the distinction between primary and secondary iconcicity is more of a cline, defined by the degree to which the "sign function" (i.e. knowing what an expression represents) is necessary for perceiving the similarity involved in iconcicity. From this perspective, the *bouba/kiki* phenomenon is somewhat intermediary in the cline.

Syncronically speaking, the iconicity of Japanese mimetics must be secondary rather than primary: once the child learns what concepts they express as part of the language acquisition process, some of the similarities could be perceived. It is much more difficult to answer the diachronic question: how the particular set of mimetics emerged in the first place? But it is clear that it must have involved a social, collaborative process, and not just a matter of Japanese speakers spontaneously externalizing their specific dynamic categories in speech. Furthermore, since Japanese mimetics are *conventional* expressions the motivation behind their meaning will be mediated by cultural norms and analogies to other expressions in the

---

[6] Ramachandran and Hubbard (2001: 18) note that a patient with damage of the *angular gyrus*, a cortical structure situated between the temporal, parietal and occipital lobes "showed no propensity for the bouba/kiki effect". This is intriguing since the angular gyrus is considered to play a role in cross-modal transfer, and may even be important for the comprehension of (novel) quasi-synaesthetic metaphors.

language, and hence often difficult to perceive. For example, while the *gi-tai-go* mimetic *noro-noro* ('drag oneself', 'walk slowly'), can possibly be related to the "laxness" of the nasal /n/ and the round central vowel /o/, the meaning of the *gi-jyoo-go* mimetic *noko-noko* ('nonchalantly') is less transparent, and possibly co-motivated by its analogy to *noro-noro*.

We can conclude therefore that Japanese mimetics take a somewhat intermediary place between iconic representations such as pictures and pantomimes (which can be interpreted even by virtue of primary iconcity), and fully symbolic and propositional language. While they may bear traces of their pre-representational roots, in particular by relying on the cross-modal mappings of the internal meaning space, they are clearly conventionalized linguistic representations, consisting of socially shared expressions and contents. These contents, however, appear to be more subjective, somewhat difficult to define, and very difficult for second language learners of Japanese (Ivanova 2002).

This conclusion is fully consistent with the analysis of Japanese mimetics presented by Kita (1997, 2001). Kita argues that the meaning of Japanese mimetics is (primarily) represented in an *affect-imagistic* dimension, where "language has direct contact with sensory motor and affective information" (Kita 1997: 380) and "vivid imagery of perceptual and physiological experiences" (Kita 2001: 420). In contrast, Kita advocates that the meaning of non-mimetic expressions constitutes an *analytic* dimension, including "quantifiers, logical operators, and semantic categories such as agent, patient and action" (Kita 1997: 1863). Apart from the iconicity of Japanese mimetics discussed above, Kita provides the following types of evidence for the need to evoke two different kinds of representations for mimetic and non-mimetic expressions.

(a) A mimetic such as *suta-suta* ('walk hurriedly') does not lead to redundancy when combined with a semantically overlapping non-mimetic expression such as *haya-aruki* ('walk hastily') in a single clause. In comparison, the combination of the latter and another overlapping adverbial such as *isogi-ashi* ('hurriedly') does lead to an impression of "wordiness".

(b) A clause with an (adverbial) mimetic cannot be combined with sentence negation.

(c) The production of a mimetic is highty associated with expressive intonation and spontaneous iconic gestures: 95% of the mimetics

> produced in a study where precisely synchronized with an iconic gesture, compared to only 36% of the verbs.

At the same time, Kita (2001) clarifies that this evidence concerns only so-called "adverbial mimetics", which are most common and least integrated into the grammar of the language, rather than cases when mimetics are used as verbs, nouns, and noun-modifiers. Thus there is not only close *integration* between the two dimensions, as Kita points out, but also evidence for a *cline* between them, with nominal mimetics being the most grammaticalized and "analytical".

This interpretation is furthermore supported by a recent pragmatic study of the use of mimetics in 4 different spoken registers, characterized by different levels of "emotive intensity" (Baba 2003). It was found that indeed the total use of mimetics correlated with the intensity level, and that *gi-jyoo-go* (the "psychomimes") where used only when the episode was narrated from a first-person perspective, involving the highest degree of subjectivity. At the same time, the nominal use of mimetics was most typical with the *least* emotive and most detached of the four levels.

Finally, we can link these studies to the concept of *mimesis* and *mimetic schema,* a mental representation involving bodily simulation that is prelinguistic, and argueably precedes and "grounds" language in phylogeny and ontogeny (Zlatev 2005, *this volume*). Mimetic schemas are dynamical structures of consciousness involving the body image, used in pre-linguistic thought and externalized in body movements and gestures. The meaning of adverbial Japanese mimetics would thus appear to correspond rather directly to mimetic schemas, even more so than the meanings of verbs, while those of nominal and verbal mimetics, as well as non-mimetic expressions would qualify as "post-mimetic". Thus our distinction mimetic/post-mimetic appears to correspond quite closely to Kita's distinction between the "affecto-imagistic" and the "analytic" dimension. Both dimensions are necessary and need to be integrated for effective communication, especially in literature.

Consider the following passage, taken from the novel *And Then* by Soseki Natsume:

> Turning to the head of his bed, he noticed a single camellia blossom that had fallen to the floor. He was certain he had heard it drop during the night; the sound had resounded in his ears like a rubber ball bounced off the ceiling. Although he thought this might be explained by the silence of the night, just to make sure that all was well with him, he had placed his right

> hand over his heart. Then, feeling the blood pulsating correctly at the edge of his ribs, he had fallen asleep.[7]

The meaning conveyed by this passage consists only in part of its propositional, analytic content, representing subjective and objective states-of-affairs such as the protagonist turning to the head of the bed, a flower lying on the floor, his memory of a loud sound, placing the hand over the heart etc. What we could call the "embodied meaning" (if the phrase was not overused nowadays) supplements this by the reader *identifying* with the protagonist and mimetically experiencing the situation from the protagonists' point of view. Notice that Natsume explicitly mentions four different sensory modalities: Proprioception (turning, placing the hand, blood pulsating), Vision (noticing), Hearing (droping, bouncing), Touch (feeling). Due to the cross-modal connections of the internal meaning space, there are quasi-synsaesthetic experiences as well: the smell of the camellia bloosom, and the temperature (warmth) of the blood. It is possible to associate in the affecto-imagistic dimension further. For example, for one of us (the native speaker of Japanese), mentioning the pulsating of *blood* and the *heart* brings to mind the color red. Hence the color of the falling blossom is also perceived (in imagination) as red! The chain of subjective quasi-synaesthetic experiences can run on: the color red induces a memory of a red sunset. The memory cold air at sunset stimulates the olfactory senses and that further stimulates the tactile feeling of a cold handrail… However, since these meanings are not conventional, they are bound to remain rather private associations. For example, the second (and non-Japanese) author "sees" the fallen flower as white, while a young Swedish poet (reading the English translation) is convinced that it is pink! Such indeterminacy has its advantages, as in the interpretation of poetry, but is problematic if it is essential for "sender" and "receiver" to be able to share similar experiences.

Japanese mimetics offer the advantage of being conventional, and thus less idiosyncratic, while at the same time evoking the affecto-imagistic

[7] This constitutes a literary translation, by Norma Field, of the original passage, which reads: *Makura moto wo miruto, yae no tubaki ga itirin tatami no ueni otiteiru. Daisuke ha yuube tokono nakade tashikani kono hana no ochiru oto wo kiita. Kare no miminiha, sore ga gomumari wo tenjyou-ura kara nagetuketa hodo ni hibiita. Yoru ga fukete, atari ga sizukana seikatomo omottaga, nennotame, miginote wo shinzou no ueni nosete, abara no hazureni tadashiku ataru chi no oto wo tashikamenagara nemuri ni tuita.*

demension. In a passage from another novel, *Wayfaring Diary*, Soseki Natsume uses a number of mimetic expressions describing the spatial movements of the first-person protagonist, as well as of other personae.

> A flight of stone steps brings back my memories. Some time I **strolled** around the "Five Mountains". Just like today I was **sluggishly** walking up the stairs which leads to the residence of monks in Engaku temple or somewhere else. Out of the gate appeared a monk in a yellow robe. He had a flat-crowned head. I ascended while he descended. When we came across, he asked in a sharp voice, "Where are you going?" I answered, "To see the Precincts", and stopped. "There's nothing within the precinct", the monk gave an immediate answer as he left **quickly** down the stairs.

The mimetics used in place of the highlighted parts are *guru-guru* ('strolling around'), *nosori-nosori* ('sluggishly') and *suta-suta* ('walk hurriedly') capturing vividly the contrast between the protagonist's motion, and metaphorically his state of mind, and those of the monk.

6. Summary and conclusions

The goal of this chapter has been twofold. On the one hand, we have tried to emphasize the following qualitative *differences* in human cognition:

- interaction is different from representation,
- procedural knowledge is different from declarative knowledge,
- dynamical categories are different from concepts,
- synaesthesia is different from language,
- the meaning of Japanese mimemics (the affecto-imagistic dimension) is different from the meaning of e.g. quantifiers (the analytical dimesion) etc.

We believe that such emphasis of differences is necessary because there is persistent tendency in embodiment theories, e.g. the "full embodiment" approach advocated by Nuñez (1999), to "resolve" such oppositions by ignoring the differences, and thus, in effect, reducing or eliminating the second and more "disembodied" side of the oppoitions. In our view, such an approach is simplictic and since we believe that conceptual – and in some cases even ontological – differences between different levels need to be maintained.

At the same time, our second goal has been to explore how structures and processes of pre-representational cognition such as dynamical categories, internal meanings space, and synesthesia can play a role in the "grounding" of mental representations (concepts) and language, i.e provide evolutionary and ontogentic prerequisites from the emergence of the latter. Figure 4 summarizes our general picture of the major different levels of meaning in a *pyramid of semiotic development*.

The rock-bottom of cognition is life itself, the *sine qua non* of all meaning (von Uexküll 1949; Maturana and Varela 1987; Zlatev 2003). Value systems derived from natural selection in evolution control the behavior and learning of living organisms. Using the complex systems modeling approach in Section 3 we showed how simple "artificial creatures" can form dynamical categores through sensorymotor coordination without any representational ability. At the same time it was suggested that by distingusihing perception and memory of past experience, categorization can be enhanced. Cross-modal mappings of various sorts are prevalent in the internal "space" which defines a coherent, multi-modal world for the subject. Synaesthesia may be only the tip of the iceberg of this, showing the importance of both correlating and differentiating modalities.

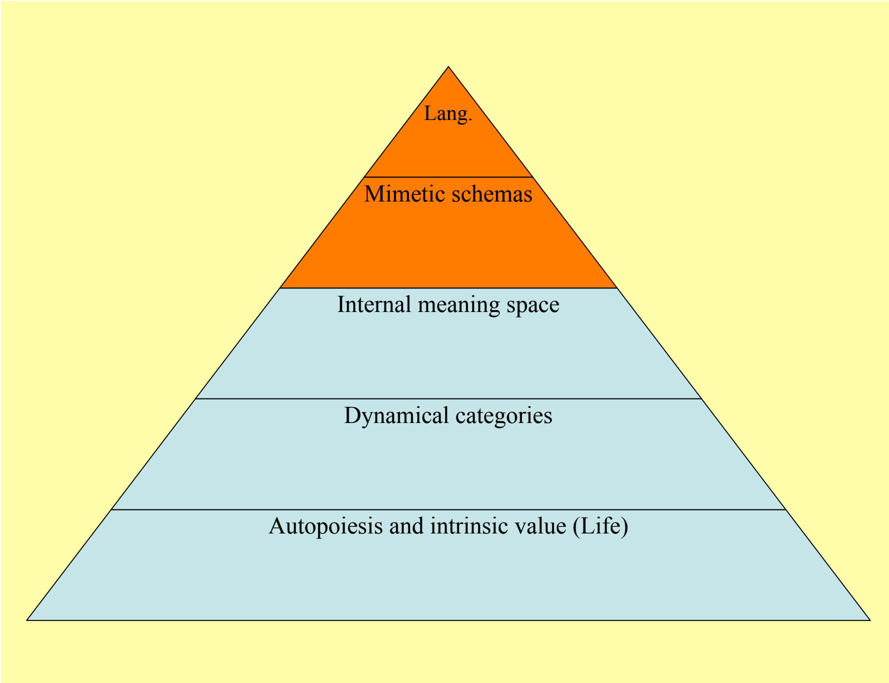

Figure 4. From Life to Language. Autopoiesis and intrinsic value preserving the identity of the organism, essential properties of life, constitute the primary basis for cognition and consciousness. The interconnectedness of dynamical categories along common quality dimensions provides an internal meaning space. Bifurcations of this space, possibly arising from self-other separation in overt or covert imitation gives rise to bodily mimesis and mimetic schmeas, which are consciously accessible mental representations. These are imagistic and affect-laden, and can be said to ground language (L) in ontogeny and phylogeny. Only the highest two layers involve mental representations.

The crucial step between pre-representational and representational cognitions occurs, we suggested, with the bifurcation of the internal meaning space into a part the focuses on the external world, and another that "looks" into memory, i.e. recall and the projected future, i.e. planning. A necessary step for this, as with any representation, is to acknowledge that a certain "expression" both corresponds to and is different from certain "content". It is possible that this breakthrough occurred precisely with bodily mimesis, making the body image the first true "signifier", but this remains so far only a conjecture. In any case, the presence and role of mimetic schemas is most clearly shown in iconic gesturing, which is universal and ubiquitous. The close synchronization of gesture and speech can be explained if mimetic schemas underly both, with iconic gestures, and certain structures of language such as mimetic expressions, being more direcltly related to the mimetic, imagistic dimension, while most of the unique properties of language are qualitatively distinct from mimesis in being symbolic, propositional and fully conventional: the top triangle of the pyramid. Of course, just like the *Cat with a Hat* of Dr. Seuss, we could cut up this top in smaller and smaller slices, with the more "abstract" ones being on the top: written language, mathematics, symbolic logic … On the other hand, actual language use does not involve the "top" alone, but the whole pyramid, and at least the imagistic dimension involved in mental simulation (imagination) and the quasi-synaestehtic associations of the internal meaning space. This is, in brief, our view of how language can be "embodied", while at the same time remaining conceptually and ontologically irreducible to sensorimotor experience.

Acknowledgements

We wish to thank the collaborators of the first author: Ryoko Uno, Hiroyuki Iizuka and Gentaro Morimoto, as well as Göran Sonesson, Tom Ziemke and Sotaro Kita for their valuable comments on an earlier draft. The first author was partially supported by a grant-in-aids from The 21st Century COE (Center of Excellence) program (Research Center for Integrated Science) of the Ministry of Education, Culture, Sports, Science and Technology, Japan, the Kayamori foundation and the ECAgent project, sponsored by the Future and Emerging Technologies program of the European Community (IST-1940). The second author was supprted by the *Language, Gesture and Pictures in Semiotic Development* project at the Faculty for Humanties and Theology at Lund University, Sweden.

References

Baba, Junko
2003 Pragmatic function of Japanese mimetics in the spoken discourse of varying emotive intensity levels. *Journal of Pragmatics* 35: 1861-1889.
Baron-Cohen, Simon
1996 Is there a normal phase of synaesthesia in development?

PSYCHE, 2(27), June 1996, http://psyche.cs.monash.edu.au/v2/psyche-2-27-baron_cohen.html
Brooks, Rodney
1999 *Cambrian Intelligence*. Cambridge, Mass.: MIT Press.
Chalmers, David
1997 *The Conscious Mind: In Search of a Fundamental Theory.* Oxford: Oxford University Press.
Cytowic, Richard
1995 Synaesthesia: Phenomenology and neuropsychology. PSYCHE, 2(10), July 1995, http://psyche.cs.monash.edu.au/v2/psyche-2-10-cytowic.html
2002 *Synesthesia: A Union of the Senses. Second edition.* Cambridge, Mass: MIT press.
Damasio, Antonio
2000 *The Feeling of What Happens. Body, Emotion and the Making of Consciousness.* New York: Harvester.
Day, Sean
1996 Synaesthesia and synaesthetic metaphors. PSYCHE, 2(32), July 1996, http://psyche.cs.monash.edu.au/v2/psyche-2-32-day.html
Dreyfus, Hubert and Stuart Dreyfus
1986 *Mind over Machine. The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press.
Edelman, Gerald
1992 *Bright Air, Brilliant Fire: On the Matter of the Mind.* London: Basic Books.
Fodor, Jerry A.
1981 *Representations.* Cambridge, Mass.: MIT Press.
Freeman, Walter.
2003 A Neurobiological Theory of Meaning in Perception. *International Journal of Bifurcation and Chaos* 13/9: 2493-2511.
Gallagher, Shaun
2005 *How the Body Shapes the Mind.* Oxford: Oxford University Press.
this volume
Gardner, Howerd
1987 *The Mind's New Science*. London: Basic Books.
Gibson, James J.,
1962 Observations on active touch. *Psychological Review* 69: 477-491.
1979 *The Ecological Approach to Visual Perception.* Boston: Houghton-Mifflin.
Gärdenfors, Peter
2000 *Conceptual Spaces. The Geometry of Thought.* Cambridge, Mass.: MIT Press.

this volume
Hamano, Shoko
1998 *The Sound-symbolic System of Japanese*. Stanford, CA: CSLI.
Hebb, Donald O.
1949 *The Organization of Behaviour*. New York: John Wiley & Sons.
Ivanova, Gergana
2001 On the relation between sound, word structure and meaning in Japanese mimetic words. Iconicity in Language. http://www.trismegistos.com/IconicityInLanguage/Articles/Ivanova.html
Iizuka, Hiroyuki and Takashi Ikegami
2004 Simulating autonomous coupling in discrimination of light frequencies. *Connection Science* 16/4: 283-299.
Johansson, Gunnar
1972 Visual perception of biological motion and a model for its analysis, *Perception and Psychophysics* 14: 201-211.
Johnson, Mark
1987 *The Body in the Mind*. Chicago: University of Chicago Press.
Johnson, Mark and George Lakoff
2002 Why cognitive linguistics requires embodied realism. *Cognitive Linguistics* 13 (3): 245-263.
Jordan, N.
1972 Is there an Achilles heel in Piaget's theorizing? *Human Development* 15: 379-382.
Kita, Sotaro
1997 Two-dimensional semantic analysis of Japanese mimetics. *Linguistics* 35: 379-415
2001 Semantic schism and interpretive integration in Japanese sentences with a mimetic: a reply to Tsujimura. *Linguistics* 39: 419-436.
Kaneko, Kunihiko and Takashi Ikegami
1998 *Evolutionary Scenario of Complex Systems Studies*. Tokyo: Asakura.
Kaneko, Kunihiko and Tsuda Ichiro
2000 *Complex Systems: Chaos and Beyond: A Constructive Approach with Applications in Life Sciences*. Berlin: Springer-Verlag.
Lakoff, George
1987 *Women, Fire and Dangerous Things: What Categories Reveal About the Mind.* Chicago: University of Chicago Press.
Mandler, Jean
2004 The Foundations of Mind: Origins of Conceptual Thought. Oxford: Oxford University Press.
Marocco, Davide and Dario Floreano
2002 Active vision and feature selection in evolutionary behavioral systems. In: Bridget Hallam, Dario Floreano, Jean-Arcady

Hallam, Gillian Hayes and John A. Meyer (eds.) *From Animals to Animates VII: Proceedings of the 7th International Conference on Simulation of Adaptive Behavior*, 247-255. Cambridge, Mass: MIT Press.

Martin, Samuel, E.
1975 *A Reference Grammar of Japanese.* New Haven: Yale University Press.

Maturana, Humberto and Francisco Varela
1987 *Tree of Knowledge: The Biological roots of Human Understanding*. Boston: Shambhala.

Maurer, Daphne
1993 Neonatal synesthesia: implications for the processing of speech and faces. Benedict de Boysson-Bardies (ed.) *Developmental Neurocognition: Speech and face processing in the first year of life.* Dordrecht: Kluwer.

Meltzoff, Andrew and Michael Moore
1977 Imitation of facial and manual gestures by human neonates. *Science* 198: 75-78.
1995 Infants' understanding of people and things: From body imitation to folk psychology, José Bermúdez, Naomi Eilan, and Anthony Marcel (eds.), *The Body and the Self,* 43-70. Cambridge, Mass.: MIT Press.

Morimoto, Gentaro and Takashi Ikegami,
2004 Evolution of dynamic categorization, In: Jordan Pollack, Mark Bedau, Ohil Husbands, Takashi Ikegami and Richard A. Watson (eds.) *Proceeding of Artificial Life XI.* Cambridge, Mass.: MIT press (to appear).

Myowa-Yanakoshi, Masako, Tomonaga, Masaki, Tanaka, Masayuki and Tetsuru Matsuzawa
2004 Imitation in neonatal chimpanzees (Pan troglodytes), *Developmental Science* 7 (4): 437-442.

Nolfi, Stefano and Davide Marocco
2002 Active perception: A sensorimotor cccount of object categorization. In: Bridget Hallam, Dario Floreano, Jean-Arcady Hallam, Gillian Hayes and John A. Meyer (eds.) *From Animals to Animates VII: Proceedings of the 7th International Conference on Simulation of Adaptive Behavior*, 266-271. Cambridge, Mass: MIT Press.

Nuñez, Rafael
1999 Could the future taste purple? Reclaiming mind, body and cognition. *Journal of Consciousness Studies*, 6 (11-12): XX-XX.

O'Regan, J. Kevin. and Alva Noë
2001 A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences*, 24: 939-1011.

Osgood, Charles E.
1959 The cross-cultural generality of visual-verbal synaesthetic tendencies. *Behavioral Science,* 5: 146-169.
Piaget, Jean
1945 *La formation du symbole chez l'enfant*, Neuchátel-Paris: Delachaux et Niestlé; English translation: G. Gattegno and F. M. Hodgson. *Play, Dreams, and Imitation in Childhood*. New York: Norton, 1962.
Pheifer, Rolf and Christian Scheier
2001 *Understanding Intelligence*. Cambridge, Mass: MIT press.
Port, Robert and Timothy Van Gelder
1998 Mind as Motion: Explorations in the Dynamics of Cognition. Reprint Edition. Cambridge, Mass.: MIT Press.
Ramachandran Vilayanur S. and Hubbard Edward M.
2001 Synaesthesia: A window into perception, thought and language. *Journal of Consciousness Studies* 8 (12): 3-32.
Rizzolatti, Giacomo and Michael Arbib
1998 Language within our grasp. *Trends in Neurosciences* 21: 188-194.
Rosch, Eleanor
1973 Natural categories. *Cognitive Psychology* 4: 328-350.
Sasaki, Masato
2000 *Chikaku ha Owaranai* [Perception Never Stops – Introduction to Affordance] Tokyo: Seido-sha.
2002 *Sentai-Shinrigaku no Kousou* [Grand Design of Affordance] Tokyo: University of Tokyo Press.
Saussure, Ferdinand de
1916 *Cours de Linguistique Générale* [Course in General Linguistics]. Paris: Payot.
Scheier, Christian and Rolf Pfeifer
1995 Classification as sensory-motor coordination: A case study on autonomous vehicles. In: Federico Moran, Alvaro Moreno, Juan J. Merelo and Pablo Chacon (eds.) *Proceedings of the 3rd European Conference on Artificial Life*, 657-667. Berlin: Springer-Verlag.
Searle, John
1992 *The Rediscovery of the Mind.* Cambridge, Mass.: MIT Press.
Sonesson, Göran
2001 From semiosis to ecology. On the theory of iconcicity and its consequences for the ontology of the Lifeworld. *VISIO* 6/2-3: 85-110.
this volume
Steels, Luc

1999 The artificial life roots of artificial intelligence. *Artificial Life* 1: 75-110.
Tani, Jun and Stefano Nolfi
1999 Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12: 1131-1141.
Thompson, Evan
2001 Empathy and consciousness, *Journal of Consciousness Studies* 8 (5-7): 1-32.
Ullmann, Stephen
1964 *Language and Style.* Oxford: Basil Blackwell.
Varela, Francisco, Evan Thompson and Eleonor Rosch
1991 *The Embodied Mind. Cognitive Science and Human Experience*. Cambridge, Mass.: MIT Press.
von Uexküll, Jakob
1982 [1940] The theory of meaning, *Semiotica* 42 (1): 25-82.
Ziemke, Tom
2001 The construction of 'reality' in the robot. *Foundations of Science* 6(1), 163-233.
2003a What's that thing called embodiment? In: Richard Alterman and David Kirsh (eds), *Proceedings of the 25th Annual Meeting of the Cognitive Science Society,* 1305-1310. Mahwah, NJ: Lawrence Erlbaum.
2003b On the role of robot simulations in embodied cognitive science. *AISB Journal* 1(4): XX-XX.
Ziemke, Tom and Noel Sharkey
2001 A stroll through the worlds of robots and animals: Applying Jakob von Uexküll's theory of meaning to adaptive robots and artificial life. *Semiotica* 134(1-4): 701-746.
Zimeke, Tom and Mikael Thieme
2002 Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response Tasks. *Adaptive Behavior* 10(3/4): 185-199.
Zlatev, Jordan
2003 Meaning = Life (+ Culture). An outline of a unified biocultural theory of meaning. *Evolution of Communication* 4 (2): 253-296.
2005 What's in a schema? Bodily mimesis and the grounding of language. In: Beate Hampe and Joe Grady (eds.), *From Perception to Meaning: Image Schemas in Cognitive Linguistics.* Berlin: Mouton de Gruyter.