CHAPTER THIRTEEN

NATURALIZING EPISTEMOLOGY FOR AUTONOMOUS SYSTEMS

JAIME GOMEZ RAMIREZ

This paper explores issues concerning naturalized epistemology and the use of isomorphism criteria in the analysis and construction of autonomous systems. A naturalized epistemic account is proposed following the constructivist paradigm. We begin by defining the basis level, where irreducible extralogical-phenomenic primitives are set out. Upon the simple primitives of basis level, further and more complex levels are defined, through a subsumption process, permitting to correlate different conceptual levels in terms of their respective primitives. Following this constructivist practice, we expect to obtain a shaped set of isomorphisms between the system form, that is, the epistemic part of the system, and a range of perceived objects and events of the environment where the system is placed.

1. Introduction

In the analytic philosophy tradition, the only reason concepts could be shared by different people was because they are disembodied and abstract. The meaning, according to this theory, would be a set of abstract relationships between words and aspects of an objective, mind-independent and external world. This extreme externalism comes from the erroneous identification of thought with language; an analysis of language cannot explain by itself the thought, and the former is an approximate vehicle to express the later. We consider that language and thought are two different ontological categories, so we cannot subsume one into the other.

In the theory of meaning we are proposing, the meaning of concepts comes through experience. Contrary to externalist and formalist approaches, words do not pick up existing entities in an objective world but express concepts which reside in the mind. The traditional sharp distinction between perception and conceptualization must be undercut.

Recent experiments have shown that the same neural mechanisms are involved in both perception and conceptualization (Gallese and Lakoff 2005, Kohler et al. 2002, Gallese 2003). We know that part of the conceptual system in the human mind arises from exercising motor schemas; so, roughly, action schemas precede conceptual though (Arbib 2006). The hypothesis we are setting out is inspired by work in cognitive linguistics (Feldman 2006), computational models (Feldman et al. 1996), and brain studies (Pulvermüller 1999). It assumes that pre-motor and parietal brain areas form a functional unity that achieves both action control and representation. Therefore, the meaning of the words is inescapably related to action and perception. Our claim is that we cannot set aside this empirical assertion. If neural mechanisms are involved in perception and action, and bodily movement plays a central role in conceptualization, it is worth building biologically plausible ontologies for autonomous systems.

Robots with a predominantly reactive behaviour may dwell in a domestic environment without crashing, successfully executing routines like avoiding obstacles or following a straight path (Brooks 1991). But to affirm that the machine possesses the concepts OBSTACLE or STRAIGHT is totally illusory, as it is the programmer, and only him or her, who has the concepts. Only autonomous systems that count on their own conceptual systems can evolve in new environments and reconfigure their goals when necessary (Sanz et al. 2010).

This paper is structured as follows. It starts by giving a definition of concepts within the context of autonomous systems. In sections 3 and 4, a critical review of classical and prototype theories of concepts is provided, with particular emphasis on the limitations built into the denotative and dyadic view of representation shared by both theories.

Section 5 establishes the basis for a naturalized theory of representation, in the light of which representation is a three term *symbol-concept-referent* relationship. The referents of the external world are grasped by the agent's mind as concepts and can be externally expressed through symbols.

Section 6 sketches the primitives or building blocks for a naturalized theory of concepts. The linguistic, sublinguistic, and neural primitives are introduced.

2. What Is a Concept and What Is It Good For in Autonomous Systems?

Concepts are the most fundamental constructs in a theory of mind. Traditionally, the study of concepts has been focused on lexical concepts, so it is worth remarking that lexical concepts are just one type of concept, while there are others such as images in our conceptual systems too. The linguistic approach to concepts is motivated by the easy accessibility and apprehension of words: we know lexemes denote concepts, but obviously this is insufficient to build a complete theory of meaning. Neural correlates and socio-cultural aspects must be incorporated into our theories of concepts[1].

We define concepts as embodied mental representations. We can use and share concepts by means of words, but the causal relation, that is, where and how concepts arise, is not only a linguistic affair, but a problem that needs to be addressed at multiple levels, with the neuronal level being particularly important.

This is not to say that concepts are simply triggered by neural structures. The deterministic explanation of how and where this is achieved is still to be constructed, but evidence suggests that perception and action circuits are both the neural basis for word meaning (Bergen et al. 2004). So, we might conclude that concepts are in part caused by sensorimotor systems in our brain.

To clarify this idea, it is useful to consider Frege's semantic theory as it is opposed to our claim that concepts are neural structures that make use of the sensorimotor systems of our brain. In Frege's "On Sense and Reference", mental representations are automatically set aside from any semantic theory on the grounds that they are intrinsically subjective – "two people are not prevented from grasping the same sense" but they can share the same sense because senses are external to our minds (Frege 1960).

Frege is confounding, however, mental representation (in his words "the very same representation") with the very same neural correlate. From the fact that an olive tree is objective, and the mental representation OLIVE TREE is subjective, we cannot infer that two different people are unable to share the same concept or mental representation. Certainly, mental representations are subjective. They belong to the thinking subject but this does not preclude them from being shareable. Indeed, different

[1] A sort of mystical or Adamic concept of language as an agent that affects matter is described in (Gallese and Lakoff 2005).

people can have the same concept even though the concepts come from different people with different brains and cultural backgrounds.

The way the brain conceptualizes is multimodal. Conceptual representation arises not only from cortex areas related to abstract thinking but also from sensorimotor areas of the brain in charge of transmitting movement signals to the muscles. It is through action and perception that we construct abstract representation of referents (external objects in the world).

3. Classical Theory of Concepts

The theory of concepts presented in this section has been predominant in the theory of mind until the 1970s, the moment at which psychological evidence began to reveal its weaknesses.

Classical Theory (henceforth CT) states that concepts are mental structures consisting of a set of sufficient and necessary conditions to be satisfied. Formally, an instance *i* falls under a concept *C,* if and only if the instance possesses the features *f* of the concept.

$$\forall f \in C, f(i) = true$$

Let us consider an example. In accordance with the compositional semantics principle that this theory follows, the concept WIDOW is defined in terms of the juxtaposition of other concepts such as WOMAN, ADULT, WAS MARRIED and HUSBAND DIED. So, an instance *i* falls under the concept WIDOW iff the next predicate is satisfied:

$$WOMAN(i) \wedge ADULT(i) \wedge WAS\_MARRIED(i) \wedge HUSBAND\_DIED(i)$$

The strength of this theory becomes its weakness from a naturalized point of view. The notion of membership to a concept is clear cut and discrete, that is to say, we can conclude that an object falls under a concept by means of checking the necessary and sufficient conditions. Proceeding in this fashion we are assuming implicitly that every instance of a concept belongs to it to the same degree. But psychological experiments have shown that not all the instances have equal footing in the concept they belong to. So, when asked to give an example of the concept fruit, apple is more frequently cited than pomegranate (Rosch 1973).

Another criticism of this view is that CT has an excessively descriptivist bias. In practice, we do not need to elaborate long lists of necessary and sufficient conditions that need to be satisfied to conclude whether an object belongs or not to a concept. From a naturalistic point of view, categorization seems to be done more efficiently by endorsing an extensional approach (rather than only the intensional approach taken by CT) which consists in contrasting the object to be included into a concept with other members of that concept. In summary, CT assumes wrongly that concepts have deterministic extensions and it does so by omitting a fundamental fact: conceptual boundaries are in most cases ill-defined.

It is important to acknowledge that an individual statement is never confirmed in isolation but instantiated in a global theory; indeed statements that in an instant t_1 are true in a successive instant t_2 can be false (Quine 1990). From this it follows that an account for temporality and cultural changes must be included in a naturalized theory of meaning. Concepts like MARRIAGE or VIRUS have been modified in our conceptual systems not long ago and they are still evolving. In some countries, people of the same sex can get married and since the dawn of the PC era a virus can be a computer program.

4. Prototype Theory of Concepts

Prototype Theory (PT) follows Wittgenstein's motto that formal criteria are neither logical nor psychological necessities. As Wittgenstein pointed out with the concept GAME, there is no unique common list of properties that games must satisfy to be considered members of GAME. It is important to note that PT does not deny that different items inside a concept might share similarities but rather holds that the similarity function is statistical and continuous and not discrete and digital as CT asserts. It is worth noting that when we talk about essential properties we are not claiming that things have essences in the Aristotelian sense. It happens that when we represent things we do so as if they had essences; this could be explained in terms of evolutionary theory since essential properties could have been the most efficient representational mechanism for survival.

Before passing to the criticism of prototype theory, it ought to be said that while in PT concepts encode properties following a similar principle as CT does; the difference resides in the fuzziness of the membership predicate. In CT, the instances are or are not members of a concept without any distinction of degree, while in PT, the items inside a category have unequal status, i.e., a robin is more prototypical of bird than a penguin.

However, experiments have revealed (Armstrong et al. 1983) that people tend to think, even for abstract mathematical concepts like EVEN, that some members of the class are more representative than others. For example, 8 is considered a better representative of concept EVEN than 48. Although this constitutes evidence for the fact that human categorization is based on prototypes it is absurd from a logical point of view.

Relativity in concept membership can lead to contradictory or unsound outcomes; because of the statistical bias of PT it could happen that one entity, satisfying some properties of a concept, could be considered a member of the concept without actually being it. To avoid this limitation, essential attributes must be provided but without reverting to classical theory. A bat could be considered a member of the category BIRD; it flies, has wings, is small etc. Thus a bat possesses a number of properties of birds, but the point is that these properties are superficial or contingent. We know a bat is not a bird because bats do not have the right DNA which shows that bats are a not birds but mammals – the evolutionary history being the key consideration.

Accordingly, using this distinction between essential and superficial properties in concepts, Komatsu (1992) has proposed a similarity space for concepts called *placeholder*, which is where the essential properties of a concept would reside, i.e., the attributes to be necessarily possessed by any entity belonging to a concept. Two remarks follow: PT using *placeholder* space might be considered a comeback to CT because it picks out properties which must be satisfied; and there is no procedure to determine the essential properties of a given concept, but, even if we could find an identifiable list of essential properties, it would not remain unchanged over time.

Our claim is that a distinction must be made between natural and nominal concepts.[2] The former would have essential properties to be satisfied, for example the DNA in plant or animal species, and the latter would be present in concepts that lack essential attributes; examples of which are easy to find in concepts concerning social conventions like KING, PHARMACIST or HOUSEWIFE.

The main problem with PT is compositionality. As Fodor has noticed, some complex concepts have no prototypes and, when they do, these do not function as their constituents. This is because the intersection operator for concepts does not work as it does in the classical set theory. In fact, a good instance of A-and-B could be a poor instance of A. Thus, for example, a good representative of the complex concept PET FISH is a golden, tiny fish inside a water tank, but on the other hand, it is a bad representative of the constituent concept FISH, which is thought of as

[2] Using Fodor's terminology, but with a different interpretation.

medium-sized and gray rather than golden or tiny. As will be explained later, an attempt to place PT inside an epistemic framework that makes it possible to set out the relationships among attributes that form a concept and the different concepts cannot be neglected anymore in the task of building a theory of concepts that could help us build autonomous systems provided with deliberative reasoning.

5. The Triadic Relation of Representation

Concepts cannot be understood in isolation but in terms of relations with other concepts inside a formal theory (Rips 1995). Accordingly, we need to arrange concepts in a systematic way, forming sets of beliefs that project natural causality (Keil 1989). Categorization is unlikely to be only based on judgements of strict similarity or typicality; concepts must instead be placed in the framework of a conceptual system.

CT and PT treat categorization as a function of similarity: an item is placed in a category if is similar enough to other category members, with a set of properties needing to be checked to ascertain whether this is the case.

The view defended in this paper contrasts with an approach that assumes a denotative view of representation according to which a mental concept is denoted by a symbol, typically a word. As it is argued below, a symbol is not a representation in itself, but is always a model for some cognitive agent, in some context.

Figure 1 shows the process of perception as a triadic symbol-concept-referent relationship. The external world referents are grasped by the agent's mind as concepts and can be externally expressed through symbols. This is because the world is populated by material things which undergo processes which emit energy to be captured and encoded by sensors (1). The sensory stimuli captured by the agent are objective and quantifiable. The properties of the perceived object can be measured; of course, the agent has perceptual limitations about what can and cannot be perceived, based on the nature of its sensors and the way they are attuned.

The patterns are instantiations of the concept's properties for certain kinds of perceptions (2) that try to match up with the encoded information in the sensory channels (3). When this computation succeeds, the referent is incorporated into the concept ontology. In other words, the salient features or properties of the referent are identified and related to the agent's ontology of concepts. The conceptual component of a sign is depicted in (4). In fact, it is an ontology of concepts which represent things or processes with common properties. According to this, the ontology of concepts is nomologically related thanks to the relationships among the properties of the concepts. Due to the lawfulness of the concept relations, learning is possible; if the cognitive agent lacked an ontology of concepts, it would have scattered options to survive.

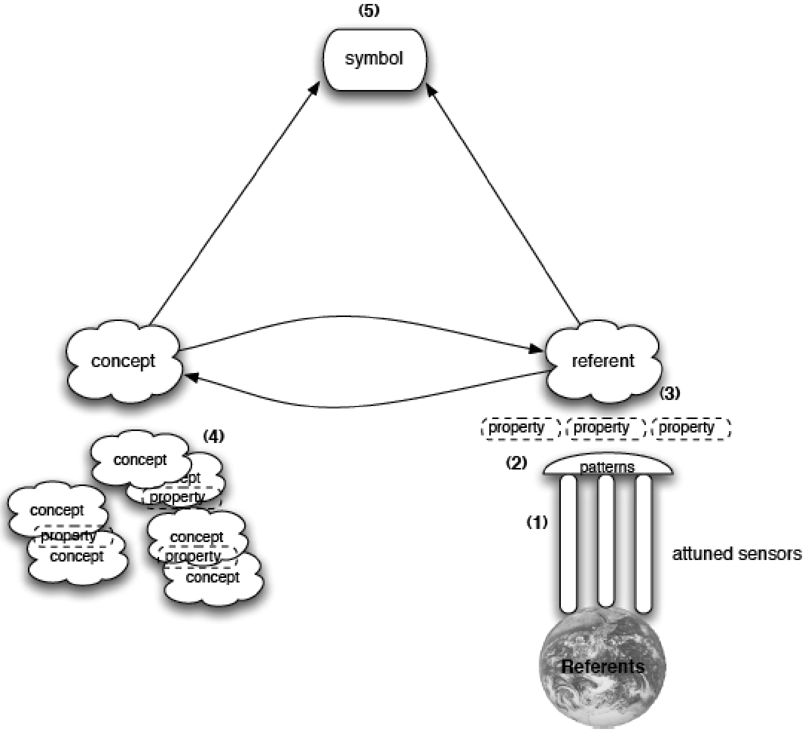

Figure 1. The triadic representation conveys to us a new definition of representation as the process of construction and relationship of signs within a system of signs.

Alternatively, if the agent, as is the case in humans, has a language or some other sign-denotative system of symbols, the relation between the external referent and the ontology of concepts can be bypassed with a symbol. The symbol (5) serves as a vehicle to share concepts within a community of agents (Gómez et al. 2008). However, there are other symbols that are not merely denotative. Instead, they permit us to infer, validate and even create novel knowledge. We call such symbols, models.

How do we know what is and what is not a model? The search for the essential features that make something a model seems a futile exercise; as a matter of fact, we make something a model by determining to use it as such (Teller 2001).

6. The Primitives in a Naturalized Theory of Concepts

The view that is it necessary to have a set of lexical primitives that, when combined, permit the construction of complex sentences is shared by almost everyone. A major exception is Fodor (1998) who considers concepts as monadic and argues that since concepts have no parts which can be separated, talk of lexical primitives would be meaningless.

Fodor claims that all concepts are innate, and arrives at this conclusion because, as he points out, we cannot explain the meaning of a word just by the means of the mere combination of other words, since that would lead to circular and ungrounded interpretation of meaning. But even assuming the correctness of this claim we do not see how concepts like BRUSH, GEAR or TAP can be innate as Fodor asserts.

The present paper avoids the innateness or the irreducibility of concepts assumed by Fodor. The problem of primitives must be reconsidered in a more thorough and multidisciplinary fashion. If we pretend to elucidate the existing connection between words and the neural configurations that constitute the cellular substrate of the mental images evoked by the words, then a theory must be constructed describing the primitives and the links between the different levels, from lexical to neural. The first step is to explain three kinds of primitives that should exist in a naturalized theory of concepts.

Linguistic Primitives

We begin by exploring the primitives at the most external level of cognition, the linguistic. Goldberg (1996), on the basis of de Saussure's works, defines constructions as the basic unit of language representation that serves to link form (phonological schemas) with meaning (conceptual schemas). Constructions are also called lexical units that may be analyzed as tuples of form and meaning, *C* = <*form, meaning*>, where *form* is the lexical expression in some of its variants (phonological, etymological etc.) and *meaning* is a conceptual schema. The two components, form and meaning, are always present, but this does not exclude the existence of additional components. For example, when the construction inherits constraints from another construction, a subcase component is included in the construction, *C* = <*form, meaning, inheritance*>.

It is important to notice two things. Firstly, form is not a single fixed word but a cluster of words that makes it possible to deal with inflected forms *(run, ran),* multi-word expression *(run off, run after)* and polysemy *(telly, television)*. Secondly, meaning may vary its structure depending on the complexity of the meaning. We will see this more clearly with the example that closes this section.

It follows, therefore, that a word is understood by means of a structured scenario which functions as the mental image necessary to ground the meaning of the construction's form. Accordingly, lexical items draw on rich conceptual structures. The theory proposed here states that language understanding implies the activation of both perceptual and motor schemas in the neural substrate. Here arises a difficulty; words are discrete but the perceptual motor schemas evoked by the words are continuous and modal.

We propose to fill this gap by putting forward a theory positing a process with two main steps: analysis and simulation. Thus, given an utterance, the analysis process determines the set of constructions that are evoked by the words of the utterance. The resulting constructions of this step serve as the semantic specification, necessary to trigger the second step, the simulation process, which consists in the execution of embodied conceptual structures. Recall simulation is an automatic reactivation of sensorimotor brain areas during concept processing.

Sublinguistic-Supracortical Primitives

To keep notation minimal, the universals in our theory are called primitives of the sublinguistic-supracortical level and configure the ontological categorization. Thanks to primitives, we count on categories that represent a deep background where the objects of a particular domain are instantiated. In our naturalized theory of concepts, nominalism is rejected, thus a symbol (/house/) refers to the external object (house) via the concept (HOUSE); otherwise, the symbol would be directly connected with the referent and we would be assuming wrongly some sort of a priori or innatist relation between symbols (words) and referents (objects). The way an agent, biological or not, is able to manipulate meaningful symbols is by sharing the ontological commitments suggested by the neural primitives described next.

DOLCE (Descriptive Ontology for Linguistic Cognitive Engineering) is a foundational ontology originally created for semantic web purposes (WonderWeb) (Masolo et al. 2009). DOLCE can also be seen as a formal ontology that captures the categories that lie behind human cognition. DOLCE is an ontology of particulars, but universals do appear in an ontology of this sort. It goes without saying that an ontology cannot be a mere classification that express facts and rules, written in a more or less formal language, but a toolset that aims to catch entities in order to place them in conceptual categories, built with a cognitive bias and forming a conceptually sound framework.

Using DOLCE we are assuming the claim that categories captured by this ontology are not related to intrinsic nature of the world (if there existed such a thing) but to cognitive items based on human perception, cultural imprints and social conventions. DOLCE taxonomy is shown in Figure 2. Four basic categories are defined in DOLCE: Endurant, Perdurant, Quality and Abstract. For our purposes we do not consider the Abstract category.

- Endurant: entities that are in time
- Perdurant: events, happen in time
- Quality: the basic entities we can perceive or measure

Neural Primitives

In the two previous sections, we have described the primitives for the linguistic and sublinguistic levels. In this section, we will set the basis for a model that correlates stimuli and neural configurations. The details of the physiological aspects of the brain are not considered here, but neurophysiological support in terms of empirical experiments can be found although it is outside the scope of this article. We take it for granted that meaning is the function of neural patterns of activation. Obviously, in order to find the neural correlates or minimum neural configuration carrying meaning, it is necessary to conceptualize perception forms, a complicated task if we consider that most of the time that we create and manipulate concepts we do so unconsciously. For example, everyday actions may be seen as routine and mechanical ones to be executed by zombie agents, that is, without access to consciousness. Consciousness only results from the action and without direct access to the action process.

Literature is rich in theories of mental states based on linguistics that tend to obviate the psychological and biological aspects of cognition. A naturalized theory of concepts, however, has to assume two epistemic positions: language is not strictly necessary to acquire and manage concepts; and if we tried to explain conceptualization only in terms of words, we would be ignoring that every single concept we use is mediated by internal physical objects, the neurons.

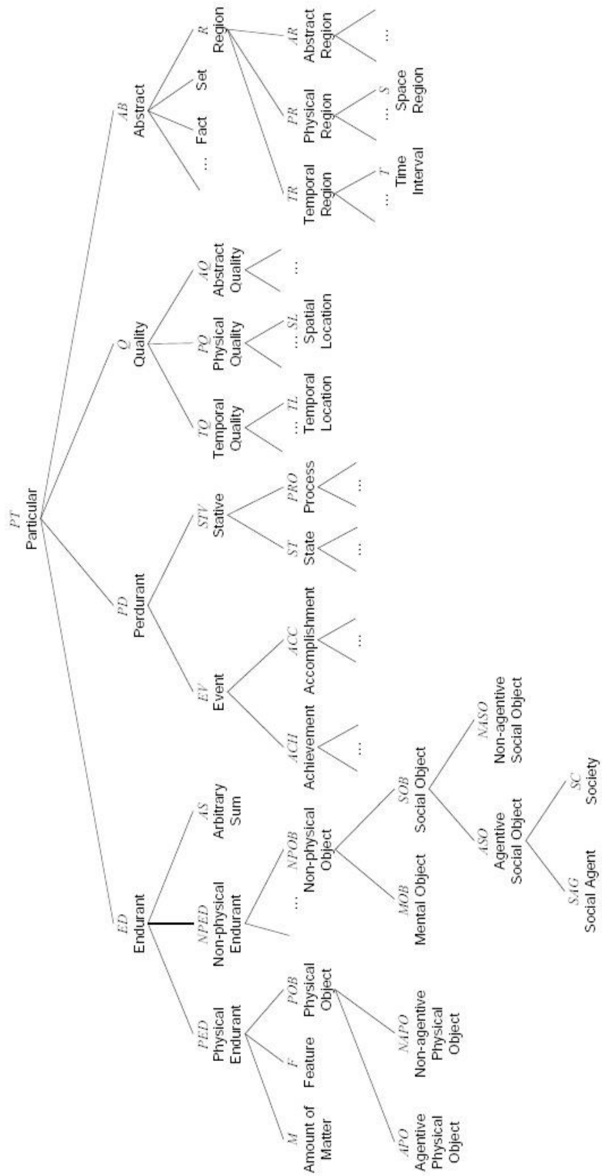

Figure 2. DOLCE (Descriptive Ontology for Linguistic Cognitive Engineering).

The motivation for the naturalized theory of concepts proposed here is to avoid circular explanation of meaning in terms of other elements such as words by outlining a model of the minimal functional unity that can carry content, i.e., neural primitives. But we have still not said anything about how our theory endows the primitives with meaning. So we need to consider the precise representational capabilities of neurons and assemblies of neurons.

Neuronal systems and content are not independent – meaning depends on neuron correlates. Cognitive neuroscientists construct statistical models that try to relate the neural response (output) to given external stimuli (input). Their experiments follow a third person approach: the input is deterministic, as are the objects shown to the subject. Nevertheless, the output is probabilistic, the potential field generated by neurons in response to presented stimuli. With this approach we obtain $P(r|s)$, which is the distribution function of the response r to a given stimuli s. In other words, we determine the way neurons spikes are generated or, in practical terms, how an internal state in the brain is caused by external stimuli.

In order to identify the neural primitives that carry meaning we need to not only know $P(r|s)$, but also must possess a model for $P(s|r)$ or how stimuli are inferred from the neural response. If $P(r|s)$ and $P(s|r)$ are known then $P(r, s)$, the probability that stimuli s and response r occur together, is easily obtained. $P(r, s)$ captures all there is to know about the probabilistic relation between r and s. Therefore, an external object and its neural correlate are linked by the highest statistical dependence.

One question immediately arises that we have to address to empirical researchers: how are the primitives for a concept to be determined once we know its neural correlate? In other words, can we find something like a principle of compositionality? It is helpful here to remember that subjective perception and reasoning are of course correlated with neural states. Different neural states can evoke the same concept, but the opposite is false; two different concepts cannot come from an identical neural configuration.

It is still unclear how the variables that constitute the neural correlates of cognitive functions, such as perception, memory, language, or consciousness, must be chosen. One of the most promising strategies for the identification of neural correlates of cognition is the state space approach originating from the analysis of dynamical systems. Neural correlates of mental states are points in the phase space of the nervous system that are associated with mental states. We can avoid the hard problem of neural properties and phenomenological states (Chalmers 1995) if we focus on the identification and isolation of these points in the phase space in order to discover under which conditions cognitive states arise and evolve in time (Fell 2004). Another interesting mathematically-based approach has been developed by the author in (Gómez 2010). There, Category Theory is proposed as a sophisticated toolkit for mental theories.

7. Conclusion and Future Work

A naturalized theory of concepts has been introduced here. Nevertheless, the theory described here is incomplete, and further effort must be expended on two questions: first, how can the basic neural primitives ascribe meaning in order to obtain the causal/computational isomorphism between neural content and external objects; and second, how to define the necessary ontological commitments, in a naturalized ontology, that describe an epistemic framework of concepts and their relationships.

References

Arbib, Michael, ed. 2006. *Action to Language via the Mirror Neuron System*. Cambridge; New York: Cambridge University Press.

Armstrong, S.L., L.R. Gleitman, and H. Gleitman. 1983. "What Some Concepts Might Not Be." *Cognition* 13 (3) (May): 263–308.

Bergen, B., N. Chang, and S. Narayan. 2004. "Simulated Action in an Embodied Construction Grammar." In *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*, 108–113.

Brooks, Rodney A. 1991. "Intelligence Without Representation." *Artificial Intelligence* 47 (October): 139–159.

Chalmers, D.J. 1995. "Facing up to the Problem of Consciousness." *Journal of Consciousness Studies* 2 (3): 200–219.

Feldman, Jerome A. 2006. *From Molecule to Metaphor. A Neural Theory of Language*. Cambridge, Mass.: MIT Press.

Feldman, Jerome, George Lakoff, David Bailey, Srini Narayanan, Terry Regier, and Andreas Stolcke. 1996. "L_0-The First Five Years of an Automated Language Acquisition Project." *Artificial Intelligence Review* 10 (1-2): 103–129.

Fell, Juergen. 2004. "Identifying Neural Correlates of Consciousness: The State Space Approach." *Consciousness and Cognition* 13 (4): 709–29.

Fodor, Jerry. 1998. *Concepts. Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Frege, Gottlob. 1960. "On Sense and Reference." In *Translations from the Philosophical Writings of Gottlob Frege*, ed. P. Geach and M. Black, 2nd Edition: 56–78. Oxford: Blackwell.

Gallese, Vittorio, and George Lakoff. 2005. "The Brain's Concepts: The Role of the Sensory-motor System in Conceptual Knowledge." *Cognitive Neuropsychology* 22 (3): 455–79.

Gallese, Vittorio. 2003. "A Neuroscientific Grasp of Concepts: From Control to Representation." *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 358 (1435): 1231–40.

Goldberg, A. E. 1996 "Making One's Way Through the Data." In Grammatical constructions: their form and meaning, ed. Masayoshi Shibatani, Sandra A. Thompson. Oxford: Oxford University Press.

Gómez, J, R Sanz, and C Hernández. 2008. "Cognitive Ontologies: Mapping Structure and Function of the Brain from a Systemic View." In *AAAI 2008 Fall Symposium on Biologically Inspired Cognitive Architectures*.

Gómez, J. A. 2010. *New Foundation for Representation in Cognitive and Brain Science: Category Theory and the Hippocampus*. PhD thesis.

Keil, Frank. 1989. *Concepts, Kinds, and Cognitive Development.* Cambridge Mass.: MIT Press.

Kohler, Evelyne, Christian Keysers, M Alessandra Umiltà, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzolatti. 2002. "Hearing Sounds, Understanding Actions: Action Representation in Mirror Neurons." *Science (New York, N.Y.)* 297 (5582): 846–8.

Komatsu, Lloyd K. 1992. "Recent Views of Conceptual Structure." *Psychological Bulletin* 112 (3): 500–526.

Masolo, C., S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. 2009. "WonderWeb Deliverable D18, Ontology Library (final)." In *The WonderWeb Library of Foundational Ontologies and the DOLCE ontology*.

Pulvermüller, Friedermann. 1999. "How and Where Are Words Represented and Processed in the Brain?" In *Language, Brain, and Verbal Behavior: Neurobiological Aspects of Linguistic Capacities and Language Processing*, ed. J.A. Argenter, 1:95–110. Institut d'Estudis Catalans.

Quine, W. 1990. *Pursuit of Truth*. Cambridge Mass.: Harvard University Press.

Rips, Lance J. 1995. "The Current Status of Research on Concept Combination." *Mind & Language* 10 (1-2): 72–104.

Rosch, Eleanor H. 1973. "Natural Categories." *Cognitive Psychology* 4 (3): 328–350.

Sanz, Ricardo, Carlos Hernández, Jaime Gómez, and Adolfo Hernando. 2010. “A Functional Approach to Emotion in Autonomous Systems.” *Advances in Experimental Medicine and Biology* 657: 249–65.

Teller, Paul. 2001. “Twilight of the Perfect Model Model.” *Erkenntnis* 55: 393–415.