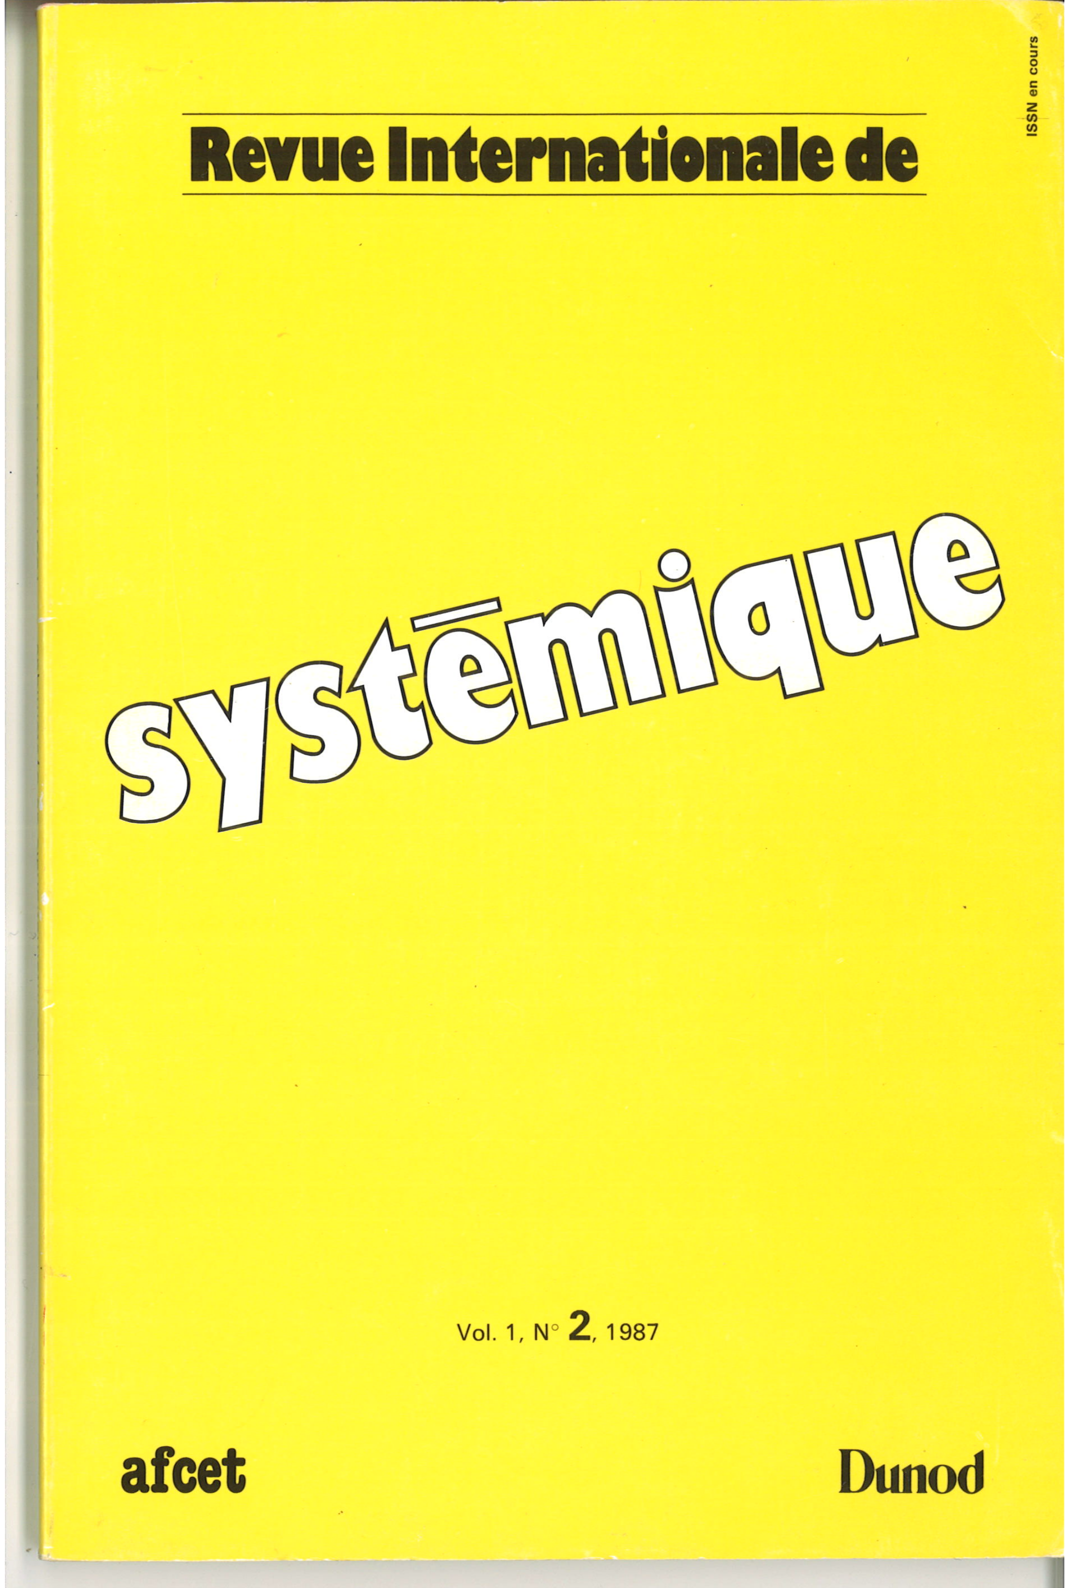


AFSCET

Revue Internationale de systémique

# Revue Internationale de Sytémique

volume 01, numéro 2, pages 209 - 223, 1987

## Experimental epistemology : background and future

Francisco J. Varela

Numérisation Afscet, décembre 2015.

REVUE INTERNATIONALE DE SYSTEMIQUE
Vol. 1, N° 2, pp. 209 à 223

# EXPERIMENTAL EPISTEMOLOGY : BACKGROUND AND FUTURE

Francisco J. VARELA

CREA - Ecole Polytechnique [1]
et Université du Chili

Résumé
Le but de ce texte est de donner un projet de programme de recherche dont l'objectif principal est compréhension nouvelle des mécanismes fondamentaux de la cognition. On insiste sur les racines conceptuelles à la base de cette approche.

Abstract
The purpose of this paper is to give an outline of a research programme where objective is an alternative understanding of the fundamental mechanisms of cognition. The main emphasis is on the conceptuel roots of this approach.

## I. *Introduction*

The main purpose of the present paper is programmatic. That is to say, its core is an outline for a research programme which projects well into the future. The main objective of this research is to intensively expand on an alternative understanding of the fundamental mechanisms of cognition, how they operate in living systems, and how can they be manufactured by man, ultimately giving the possibility of truly intelligent artifacts. The main emphasis of this paper is in situating this way of approaching the cognitive issues in its conceptual roots.

Thus, I have written this with the assumption that the reader already

1. 1, rue Descartes, 75005 Paris, France.

has some dissatisfaction with the current or mainstream research programme in the cognitive sciences and artificial intelligence, and that he is actively engaged in an examination of alternatives. Bluntly stated, I do not think that the alternatives are forthcoming unless we are prepared to probe deeply into the roots of our scientific discourse altogether. It is only at a deep foundational level that alternatives can even be seen as alternatives. Thus, the name *experimental epistelology*, originally coined by Warren McCulloch [1], is used here to denote this orientation, for obvious reasons.

We present the conceptual background of such an alternative viewpoint in the first part of this proposal. This presentation is based on over 15 years of research and results in brain and cognitive sciences in a few laboratories including my own. From brevity's sake, I refer to this orientation as the Santiago school for reasons that will be made clear below. It is not just a tentative programme but one that has proven to be consitent, workable, and capable of generating concrete alternative explanations and experimental approaches for various concrete problems. Having secured an alternative programme well grounded in the sciences, it is possible to move this line of thinking into the next interesting steps. Among these interesting steps are the design of man-made cognitive devices, and the bridges to autonomous phenomena in the social realm.

The strategy I have chosen here is to reflect on the foundational ground on which current cognitive sciences stand, so that alternatives can be seen as such. I will describe the proposed alternative linking it with the historical watershed of the Wiener/von Neumann views at the very inception of cybernetics, and how some issues discussed at the time became submerged under a dominant information-processing approach still in vogue. Hopefully, by the end the reader will have the ground, its connection to the history, and the scope of application, all in one piece. Please observe that my intention here is not to do history of science, but rather to reflect on an orientation for research and its historical lineage.

## 2. *Background*

### 2-1. We live within a tradition

Answers to the question what is cognition and how to make explicit its mechanisms, are going to be a dominant theme in the next decades, and its consequences (or lack of them) by far the most significant at the technological level. This interest is shared by many beyond the academic world, and the understanding of what such a question presupposition is by far the most influential in terms of the policies towards finding what are deemed adequate answers.



Just a few years ago there was no household meaning of terms such as 'word processing', 'information management', or 'knowledge representation'. In discussing the future of the cognitive sciences and artificial intelligence, such terms are taken as granted, and any discussion about them assumes already some common background. In other words such discussions exists within a given tradition, they are historical and not mere 'objective' facts. The tradition shared by western-science is ever present in our daily 'common sense'. Our orientation here is to open new possibilities at the level of this assumed, but usually invisible, tradition. Evidently this document cannot be such a re-orientation, but it can point far enough to entice the reader into its exploration [2].

### 2-2. Basic qualities of our tradition

For the present purpose let's agree to call our dominant tradition *representationist*. Its a token name for a multifarious network of assumptions since the beginning of our historical roots in Greece, with only a few aborted side tracks. The main pillar of such tradition is simply stated : objectivism. That is to say, the basic notion that knowledge is acquired by obtaining an adequate representation of the true laws and properties of the world, and science's task is do so through successive validations of hypothesis. So pervasive are these notions about correspondence, representation, and knowledge, that raising any discussion about them is immediately termed 'philosophical.'

Certainly in the cognitive sciences at large, such pre-suppositions are dominant, as is currently called the *information-processing approach* whether it appears in neurosciences, psychology, linguistic, artificial intelligence, or management. Simon summarizes this orientation well :

«1. Cognitive systems are symbol systems, achieving their intelligence by symbolizing their internal situation and states through symbol manipulation;

2. A theory of cognition is to find the appropiate rules or programs of symbolic manipulation so that when the system operates in the appropiate environment, it produces the expected behavior [3].

Although this statement can be re-interpreted in the particular context of any of the disciplines mentioned above, it is perhaps clearest to see them expressed in the design of expert systems.

2-3. There are alternatives to the dominant tradition

Obviously we could not have stated what we have said so far, unless there was already a horizon within which an alternative presents itself. Here, we call this alternative orientation the *middle way* : a move that goes beyond representationism and any form of subjectivism or relativism which is its logical opposite [4].

One major stream contributing to this orientation is the continental school of hermeneutics where the act of interpretation is taken as primary (instead as a residual error). Interpretation originally meant that an understanding of a text always arises within a given contextual horizon of human life, and yields a new interpretation. There is no question of what the text'itself'says ; what is fundamental is the cycle between text and interpretation in an ever changing co-specification, or hermeneutical circle. Some thinkers [5] have taken this old hermeneutical tradition beyond the abstruse interpretation of old texts, into being the center piece of what cognition is all about. The expansion is evident : texts are substituted by any form of interaction or coupling we have with the world. Hence interpretation is central to perception, action, and language altogether.

We have followed the hermeneutical move beyond representationism, since its is very illustrative of the central point made here. The consequence of taking such a stance are dramatic. In fact, let us reread the sentence 1-2 above as follows :

(1) human beings inhabit a'real'world with objects bearing properties, and such properties do not depend on interpretation ;

(2) perception is a process whereby such objects and properties are represented in such a way that appropiate computations on them can be made for adaptive behavior.

Most people would find such statements completely transparent and reasonable. Yet, from the middle way point of view, objects and their properties and the actions we perform (i.e. our interpretations) are inseparable, in fact, they are mutually definitory. We find ourselves already immersed in the cognitive activity, thrown into it, and in no way can we step outside of this circle of action-interpretation.

This brings into question the very notion of representation : what's prevalent is the actual viability of my engagement with the world in its immediacy, not whether it is a processed picture of what really is. One of the major points here is the substitution of the need to have representation for an understanding of cognitive mechanisms. We will come back to this issues in greater detail later.



2-4. A fresh look at the question of cognition

Having taken our orientation in an alternative (i.e. non-common sensical) direction, our task is to reformulate scientific questions in a way that are consistent with this new orientation and shed fresh ligth into these question as scientific questions. Before we do so, however, we need to make clear at the outset a few fundamental ideas.

First since we have suspended the need to give objects and their properties an independent status, our considerations have to be guided by bracketing objectivity. This is to say, we cannot be guided by assuming that the objects contain their own criteria of validation. This assertion will become of central importance both in the detailed understanding of the brain process, as well as in the design of man-made artefacts.

Secondly, since we have situated ourselves squarely within a tradition which is cognizant of its interpretive activity, the only departure point is to say that we are capable of making distinction of unities (or systems or entities) in the world. Such distinctions can be iterated and thus bring forth composite units and relations between their components. In what follows it will be useful to have ready at hand the difference between organization as the necessary relations which distinguishes a composite unity, and structure as the mode of instantiation of that unity. Organization is that of a system which is invariant (while it persists in our field of description), structure is that of a system which changes (while it persists being the same unity).

Within this alternative frame of reference, how are we to approach the understanding of cognition as a biological phenomenon ? This is the central question addressed by our research during the past 15 yeras [6]. In a nutshell the answer lies in a clear understanding that our alternative orientation brings to the fore the crucial importance of the autonomy of living systems. In fact, what replaces the search for representation as the key to cognitive mechanism, is the search for the ways in which the world for a living systems is inseparable from its own living, their autonomous self-assertion expressed through their living. Thus, cognition as a biological phenomenon takes on radically different dimensions : we have effected a shift from a cognitive question to one about the organization/structure of a class of systems. In the following Sections we unfold this central intuition.

2-5. From Turing to autonomous systems

For the purpose at hand, the easiest way to proceed is to step back from a moment from any specific domain (such as the biology of the nervous system ), and lift the question into a change of paradigm in systems theory and cybernetics altogether. This is useful, for these disciplines condense the very essence of the mechanism and principles common to all the regional disciplines that concern us here in a more compact form. We assert that the dominant paradigm in cybernetics, hereinafter called Turing cybernetics, needs to be extended into a more encompassing paradigm, hereinafter referred to as the cybernetics of autonomous systems. The relation between these two is homologous to the relation between Newtonian and relativistic mechanics.

The contrast between these two paradigms can be made at the two complementary level of (a) how a system is characterized, i.e. what constitutes its identity, and (b) how is it seen in its relationship with its surrounding medium, i.e. what constitutes cognition.

(a) Identity :

For Turing cybernetics the characterization always takes the form of input/output description mediated via a internal state dependency of processing mechanism. This is of course, what he have become accustomed to call a system altogether, and which achieved a transparent form in Turing's formulation.

For an autonomous system, however, its characterization relies fundamentally on a description of the internal interdependency of its components, i.e. their operational closure. This change of orientation can be further elucidated by stating a generalization of the many empirical observations of systems in nature and simulations studied under this optics : Identity is the maintenance of operational closure through the production of internal coherences [7].

(b) Cognition :

For Turing systems, the mode of relation of the system in its world is straight-forward, for it is given by their inputs, and thus what ought to be understood is the mechanism of representation. This is, of course, what we have become accustomed to call informational flow altogether.

For autonomous systems this is not the case, since having no inputs or outputs, they can only bring forth relevant distinctions which constitute their relevant world. Again a general principle from empirical studies can be formulated thus : Cognition is the conservation of identity through a history of uninterrupted structural coupling.

It should be observed that in this characterization of how to



address a system from our alternative orientation, objectivity is in fact bracketed, since there is no assumption of a fixed set of world properties for the system. Structure in the medium exists to be sure for us as observers, but the whole point is to let the dynamics of conservation of the autonomous system reveal what in it is relevant for itself, in a continued history of unfolding of meaning. This orientation amounts to an entirely new mode of inference about natural phenomena, for the reference point is not the properties of the world anymore, but rather the history of coupling mediated by the richness of the internal coherences of the system itself.

The homology to the Newton – Einstein transition is now clear. If we insist on holding as fixed a specific set of properties of the world, rather than permit a history of coupling to reveal it in ever changing possibilites, we arrive at a Turing type characterization. Such an approach is surely useful in some contexts. Our point is only that it seems to be very limited to address the problems raised by life and mind. For them, we need to relax the grip of a fixed world, and an inseparable co-specification of system and its world throught the internal coherences to arise.

## 2-6. A fresh look at the nervous system

The stated shift from Turing to autonomous systems, captures in abstract terms the essence of our pragmatics implied by the middle-way orientation. To reveal its full flavor it is necessary now to see it at work to produce a fresh understanding in some specific situations. I have discussed these applications extensively elsewhere for the key areas of the nervous and the immune system, and well as for evolutionary processes [7]. We cannot cover all of that ground here. The reader is referred to the cited papers for details.

## 2-7. A fresh understanding of language

We have stressed implications of our alternative orientation for the understanding of basic process of perception and behavior altogether. With that basis, it is a natural task now to expand that orientation into language also as non-representational process.

In fact, it is enought to realize that a living system as an autonomous entity can have another another living system as an autonomous entity as part of its medium wherein it maintains a history of uninterrupted identity. Thus the two (or more) living system become sources of mutual perturbations for each other, and thus they can recurrently generate a domain of stabilized interactions which an observer can describe in semantic terms. Such linguistic domains occur in multifarious ways in nature (e.g. bee dance, group role in primates, courtship sequences). A language proper arises when (and only when) the linguistic domain can make distinctions within the linguistic domain itself, i.e. when a recursive capacity appears. Our own language is of this kind, and further, in doing what we are doing right now we operate within such a dynamics [8].

One of the key changes here is that language is not an instrument for communicating information. Rather it is a web of common actions which exists only insofar as a shared history of structural coupling exists, and meaning lives only within this shared history. Just like in the case of perception, it is hopeless to attempt to situate meaning either in language itself or inside of the brain of a participant. The key phenomenon of language is the further transformations within coupling when behavior is ineffective.

In the middle way approach, language arises naturally of an autonomous view of living systems, including an understanding of meaning and its origin, and the kinds of dynamics that language can mediate. This continuity between perception and language is a powerful feature of our approach since it constitutes the link between the biological and the social realm, as discussed extensively elsewhere [8].

### 2-7. Conclusion

In the preceding pages, we have sketched the basis for an alternative framework within which to freshly ask question about knowledge, the origin of meaning, roots of knowledge in the biological world, and the continuity cognition-language both in nature and human beings. The reader should be reminded again that this document cannot be effective in displaying for him in full an alternative of our dominant tradition. The intention has been to orient him into possibilities, and on the other hand to point to enough scientific work to make it possible that such an alternative – working without representations, autonomous systems, structural coupling as the origin of meaning, and so on – is indeed a comprehensive and workable alternative to the representationist programme, and worth expanding further.

## 3. *Future*

### 3-1. Ontological research and design

The cornerstone of our approach to applying the present middle way perspective is in fact simple once we have grasped the basic points



outlined above. A man-made system or a natural system with the capacity to exhibit cognitive behavior – be it perceptual or linguistic – should satisfy the criteria outline in 2-5. for an autonomous system, having :

1. an invariant organization under structural change capable of constantly producing internal coherences as a result of its operational closure ;

2. an ongoing structural change resulting from its coupling with the medium (which could be itself) leading to an interrupted pathway of identity conservation.

In this process the history of structural coupling will inevitably specify a domain of relevant distinctions for the unit, and whenever its identity is conserved a behavior will exist for the conserver. If the structural plasticity is rich enough, such relevar distinction and behavior will become recurrent and hence a history of shared meaning will arise.

To this entire interwoven process we call hereinafter *ontological dynamics* for it puts into relevance the ontological or constitutive dimension of the unit's autonomy. Whenever such ontological design takes place in a space of components where the relevant dimensions are those of actions in space and time, perceptual behavior will arise. When such ontological design takes place between such systems, a linguistic domain will result.

I want to be very clear at this point that in this description I have made no distinctions between the research of natural system and the construction of artificial ones. In other words we are talking about ontological research *and* design. This is deeper than just an economy of words. It also stays that the construction of such devices is, in an interesting sense, the method of validation of the hypothesis.

We should also be careful to realize that we are not talking here of a clever way to understand how a system can acquire a good representation of our environment. Instead we are talking of systems capable of coupling with us, and, through coupling, of revealing to us what is of common relevance. Thus the researcher or designer is not to attempt the (impossible) task of making explicit what 'context' is, but allowing such context or background to emerge by an ongoing history of common coupling. Only within the hermeneutics of common facticity can truly interesting, i.e. creative behavior arise. The antithesis to this proposal is, for example, the design of expert systems, with the attempt to explicitely codify all of common sense into rules. The *measure of intelligence* in this sense, is not that the system imitates what we do, but rather the extent to which it can enter into a domain of shared coupling with us.

3-.2. Early Origins

At this point, the reader familiar with cybernetics and brain theory literature might be tempted to say : «Well, after all, all what you are talking about is a new version of the self-organizing theories and learning machines of the 1960's». In fact, after the War, there was a stress on cognition as arising from creative structures which could hardly come from the observer projecting his pre-understanding into a device or the brain, however sophisticated this projection was claimed to be. Consequently, there was a flurry of theories about structures which could 'evolve' of 'learn' under various conditions [9]. After a few years all such efforts were abandoned, and have not been renewed.

The key ideas of cognitives science in general, and of experimental epistemology as addressed here, have their most direct roots in the surge of activity after World War II in the USA, and easily identified with the figures of John von Neumann and Norber Wiener [10].

Further, the dozen meetings sponsored by the Macy Foundation could be said to be the concerte birth date of the whole complex of ideas with which we are debating up to this day. I shall not discuss these historical aspects here.

For our purpose, it is only necessary to remember that within a short span of time of about ten years virtually all of the ideas of cognition in animals and machines were sketched. Equally important, after this decade of great creativity only a subset of the remained on the mainstream of science, mostly around the notions put forth by John von Neumann and Alan Turing culminating in the design of the digital computer and the establishement of computer science as a discipline. The computer provided, in turn the basis for the definition of a whole range of metaphors on information-processing as is currently understood e.g. by Simon as quoted before. The philosophical core of this trend is, precisely, the view that cognition works by an astute representation of the information offered by the world.

The towering dominance of the von Neumannian paradigm of universal calculation completely obscured the exploration of the autonomous and self-organizing quality of the living, and the theories suggested by them, although many researchers (including Wiener) recognized how revolutionary these notions were. However, the experience was that such an approach was totally ineffective in connecting to empirical research, and even the most rudimentary interesting behavior in articial systems was not forthcoming. After a brief spring of popularity, this line of research was discontinued in the late 60's. Since then, the current approach to design became the dominant paradigm for AI by posing knowledge representation as the central issue, and neuroscience followed suit.



The reasons for the decline of these budding ideas are, with the benefit of 20 yeras of distance, easy to understand. First and foremost, it was an intuition that run counter a dominant tradition that had been central in science for a long time. Hence to fructify, it would have had to be aware of being a radical change in orientation and to make its alternative foundations more explicit. The time was not receptive to such foundational discussion, as is clear in the relative unpopularity of epistemologically-minded researchers such as W. McCulloch, Heinz von Foerster, Gregory Bateson, and Wiener himself. In a word, those researchers were looking in the right direction too early to be viable.

Secondly, the guiding images for the specification of the systems closure was crude, and based on the idea that simple components in vast numbers would do it. Two decades of research in neurobiology, immunology, and evolutionary studies have proven that biological autonomy is based on highly structured principles, the result of millions of years of natural drift. Some of these general principles can be lifted and adopted into ontological design, a rich source of ideas that was not available then.

Thirdly, the scale at which such autonomous system can be studied was limited due to capacity and speed of computers 20 years ago. This is, of course, valid for any strategy at that time, but specially for this one where simulations in vast scales is needed to weed out useful strategies, and millions of repetitions necessary to have (as is required by ontological design) a historicity of a cognitive domain. It is ironical to think that the current von Neumann architecture could be the womb within which alternative autonomous non-von-Neumann architectures could be designed for a future generations of truly cognitive machines.

But of all this tentative explanations, by far the most interesting to us is the first one : the difficult process of addressing foundational issues at an unreceptive time. As I said before, only a handful of researchers continued to pursue these intuitions. In the USA they were clearly centralized at the Biological Computer Laboratory in Illinois under the direction of Von Foerster. Among the most notable ones there : W. Roy Ashby, Gothrand Gunther and of course Heing von Foerster. Other scattered figures should be mentioned : Stafford Beer, Oliver Selfridge, and A. Rosenblueth.

It was through von Foerster that this set of preoccupations dovetailed with Santiago and Humberto Maturana in the 1960's. Up to that time, his direction was wholly neurobiological (including the famous article with Lettvin, McCulloch and Pitts on «What the frog's eye tells the frog's brain»), but independently nourishing a dissatisfaction with the dominant trend in neurobiology. For example, at that time he expressed his dissatisfaction by atttempting to describe living system as 'self-referential'. But the full depth of this questioning of biological foundations of cognition did not become apparent until 1969 with the publication the article «The neurophysiology of Cognition» written in Illinois for a Wenner Grenn Symposium organized by Heinz von Foerster. This paper can be said to be the birth certificate of the Santiago School. I had started as student of Maturana in 1966, and by this time became fully involved in the development of these ideas in my own way.

After 1969, a steady stream of work started to appear, the first one being the formulation of the notion of autopoiesis written jointly with Maturana in 1971, although not published until later. Since then, there have been many contributions which fall more or less under the inspiration of this initial line of research [6]. In the first part of this paper I summarized how these ideas look today from a foundational angle. My intention here is not historical, but conceptual, so I will not trace in detail the successive publication of key ideas and experiments. Instead, I shall present the main departure of the Santiago school from the tradition of experimental epistemology as it existed up to 1969 :

1. the explicit grounding of cognitive processes on the organization of the living itself and the characterization of the living through its autonomy. (Only later did this concern with autonomy and autopoiesis find resonance with self-organization theories) ;

2. the explicit critique to the theory of representation as the guiding thread for the operation of the nervous system ; the proposition of the operational closure of the nervous system as its proper mode of operation ;

3. the extension of this approach to other basic phenomena, specially evolution, learning, and language in a coherent framework.

Thus the Santiago orientation represents both a continuity as well as a departure which is just as radical.



### 3-3. Current convergent approaches

In the current and recent scientific discourses there are two recent trends which, though different and independent, have an interesting degree of convergence with the orientation of the Santiago shcool we have been discussing. Thece come from two disparate fields : (1) the re-discoverty of the issues around self-organization in physics, and (2) the re-discovery of some self-organizing procedures in AI. The first of these two derived mostly from the work of I. Prigogine, is also extensively discussed in other documents so I shall not touch upon it. However, a brief words about (2) is relevant, since closer to cognitive problems.

This trend is the so-called neo-connectionism [11]. The problem is to find solutions to difficult AI tasks (such as vision) by a strong appeal to biological prototypes. The main motivation in this line of research is the need for a robust framework for parallel computation, without which the possibility of approaching biological performances is – even within the traditional AI tradition – not feasible.

With this motivation, this line of work has produced tools for the simulation of strongly interconnected, plastic changing systems capable of achieving internal coherences through, for example, relaxation methods [12]. It should be clear that these researchers are not concerned with a shift in orientation within their tradition, but rather with finding efficient algorithms to solve the issue of parallelism. Their overall question remains one of building representations of the system's environment, be it in a linguistic or in a perceptual domain.

A parallel trend has emerged over the past few years, partly in overlap with the above AI tendency, among some researchers in neurobiology and cognitive spychology to look for robust principles capable of explaining the emergence of large scale coherences in complex neural-like nets [13]. Among the new results in this area is the fact that statistical reasoning could be used do characterize generic behavior of the network [14]. Again, these researchers are not centrally concerned with changing the question about what is cognition, but rather with effective solutions to the way neural representation could arise.

In summary then, both in AI and in cognitive neurosciences recent trends have emerged which retake some of the forgotten themes of the learning machines of the 1960's, into a new context of either searching useful parallel architectures, or for good models of neural behavior conceived as a representational machine. However, it is also equally clear that this research itself raises relevant epistemological questions which must be, willy-nilly, addressed.

### 3-4. Future : Structural drift with autonomous systems

I have been presenting a theme with several voices in parallel. It is a critical epistemological framework for the understanding of cognitive processes. It is a statement about a cybernetics that needs to be developed. It is a sketch about possible research with natural systems such as the nervous or the immune systems (where the system's identity and its coupling is definable). It is a programmatic approach for the design of artificial systems. It is an orientation which has precedents in the recent scientific history, but it also represents a radical departure from them. All of these threads of our main theme are closely interelated, and inseparable. What has been published so far with this orientation in all of these areas makes it clear that the gap between formulation and practice can be carried out.

One of the most interesting next steps to me is the last one, the design of artificial , systems, which merits a few closing words. We would like to conduct (for the first time) a full exploration of the possibility of man-made autonomous systems capable of a prolonged and rich history of structural coupling with humans, in such a way that a common structural dirft is created together wherein a shared world of meaning and action can be brought forth.

Such a statement can only be understood by bringing into play all we have said so far, and would not be possible without the emphasis of foundational examination we have pursued. It also would not be possible without the experience and some of the tools developed within the representationist tradition. It is fair to say that the current tradition – like all traditions – has its weaknesses and its strengths. The strength of the representationist tradition is pragmatic effectiveness for limited engineering tasks ; its blindness lies in its inability to understand context and creativity.

### Références and Notes

(1) W. McCulloch, *Embodiments of Mind*, MIT Press, Cambridge, 1960.

(2) This presentation owes a great deal of influence to the reinterpretation of the hermeneutical tradition, as developed by Fernando Flores. See specially F. Flores and T. Winnograd, Understanding Computers and Cognition, Ablex Press, New Jersey, (forthcoming)

(3) H. Simon, *The Sciences of the Artificial*, MIT Press, 1972, p. 80.



(4) The designation middle-way for the Aufhebung beyond the dialectical opposition between objectivism and relativism is taken from the medieval indian epistemological schools, best represented by Nagarjuna. The original word is madhyamika. See Candrakirti, The Lucid Exposition of the Middle Way, Shambala, Boulder, 1980 (trans. by M. Sprung).

(5) H.G. Gadamer, *Wahrheit und Methode*, Jcb Mohr, Tubingen, 1960.

(6) The exposition of these ideas is dispersed in various technical articles. Two foundational pieces are to be found in H. Maturana and F. Varela. Autopoiesis and Cognition, D. Reidel, Boston, 1980. A comprehensive exposition is F. Varela, Principles of Biological Autonomy, Elsevier North Holland, New Yorl, 1979. For a recent exposition see H. Maturana and F. Varela, The Tree of Knowledge, New Science Library, Boston, 1985 (in press) and F. Varela, *L'Autonomie du Vivant*, Eds du Seuil, Paris, (forthcoming).

(7) Cf. Varela, 1979, *op. cit.*

(8) Cf. Maturana and Varela, 1985, *op.cit.*

(9) N. Nilsson, *Learning Machines,* McGraw-Hill, New York, 1965.

(10) S. Heims, John von Neumann and Norbert Wiener, MIT Press, Cambridge, 1980.

(11) J.Feldman and D. Ballard, Connectionist models and their properties, *Cognitive Science*, 1982.

(12) T. Sejnowski and G. Hinton, Separating figure from ground with a Boltzmann machine, in : M. Arbib and A. Hanson, (Eds), Vision, Brain, and Cooperative Computation, MIT Press, 1985.

(13) S. Grossberg, *Studies in Brain and Mind,* D. Reidel, Boston, 1982.

(14) W. Freeman, *Mass Action in the Nervous System,* Academic Press, New York, 1975.