

The Past-Future of Cybernetics: Conversations, Von Foerster, and the BCL

Paul Pangaro, PhD
New York, NY
USA
pan@pangaro.com

ABSTRACT	2
INTRODUCTORY REMARKS	2
Intentions	2
Outline	3
ORIGINS.....	3
Jerome Y. Lettvin	3
Humberto R. Maturana.....	4
Andrew Gordon Speedie-Pask	5
Heinz von Foerster	5
Experimentalists	6
CLOSURE.....	8
Nervous System Loops	8
Biological Computer Laboratory	9
Meta-analysis via THOUGHTSHUFFLER	10
Common Concepts	12
INTERACTION WITH THE ENVIRONMENT	13
‘Looping-through’ and Autotopic Mapping.....	13
Von Foerster’s Eigen Functions vs. Strange Attractors	14
Pask’s Concept of ‘Concept’	14
Maturana’s Niche Model.....	19
Pask’s Musicolour Machine	19
INTERACTION VIA THE ENVIRONMENT	21
The Constructivist and the Physicist	21
Pask’s ‘Theory of Everything’	21
Maturana’s ‘Do You See The Flower?’	23
STILL TO EXPLORE.....	23
HAPPY BIRTHDAY, HEINZ	24
Question #1	25
Question #2	25
‘GETTING WHAT WE WANT’	26
EPILOGUE	27
ACKNOWLEDGEMENTS	27
SPEAKER BIOGRAPHY	27

ABSTRACT

To speak 'Biological Computer Laboratory' also speaks 'Heinz von Foerster.' To invoke von Foerster also invokes the BCL community that he gathered through his unerring identification of original thinkers and his unparalleled clarity about second-order cybernetics. Having chosen well his lab's collaborators, von Foerster contributed seminal thinking that became foundations and superstructures for practitioners in the generations that followed.

What contributions to cybernetics were rooted in the BCL? What insights did von Foerster himself offer, such that his collaborators could stand tall on his shoulders and see more? With the benefit of twenty-five years' hindsight, the speaker will analyze the published outcomes of the BCL and conjure a picture of von Foerster's influence on two of his collaborators, Gordon Pask and Humberto Maturana. A *post hoc* construction drawn from personal relationships with the protagonists, the talk will offer a unification of major threads of cybernetics, its concepts of memory, organizational closure, and circularity, and show how von Foerster is inextricably woven in.

INTRODUCTORY REMARKS

Ladies and Gentleman, it is a pleasure to be here in Vienna for this celebration of the life and work of Heinz von Foerster. I very much appreciate the opportunity to speak at this conference. It is a privilege to have known Heinz, to have been near him, and to have been encouraged by him.

[Referring to video running on projection screen] What you see here is a brief video showing Heinz's collection of cybernetic books, and the personal inscriptions to Heinz from each author. Warren McCulloch, Ross Ashby, Gordon Pask, Stafford Beer, Ernst von Glasersfeld, and others have all written notes to Heinz on the title page of their books, in tribute to him and to his influence on them.

Intentions

My abstract for today's talk implies that I will describe to you what the contributions of Heinz von Foerster and the BCL have been, are, and continue to be. When I wrote the abstract and submitted it, long before I sat down to write the paper, I wanted to specifically ask, What were von Foerster's insights, those that came from him and were transmitted to others, not only to the many generations of students, but also to the other extraordinary practitioners in the field, specifically Gordon Pask and Humberto Maturana? Hence, the title of the talk, 'The Past-Future of Cybernetics', a view of the past from the 'now', by us in the future. I wanted to uncover something of the **conversations** starting with von Foerster at the BCL and involving these other great collaborators. I chose this abstract because it has always been of great interest to me to try to understand the nature and scope of these inter-relationships [Pangaro, Paul (1995). *On the Shoulders of a Giant: Eigen-Stories of Uncle-Heinz*. <http://pangaro.com/abstracts/ASC-95-on-HvF.html>. Retrieved August 25, 2005.]. I wanted to draw a clear line from Heinz's work and thinking, forward in time.

However, Ladies and Gentleman, I regret to inform you that, after much review and deliberation, I must confess that the **goal** stated—to draw a line *from* Heinz von Foerster *to* these other individuals and the future of the field—is irretrievably and inevitably *impossible*. As Humberto Maturana might say: 'Plans never work.'

I hope that rather than disappointing you with that statement, I delight you with its implied reference to Heinz von Foerster. Those of you who have read Heinz's papers know that this was a consistent trick of his: To begin a paper, say at a conference on 'self-organizing systems', by declaring emphatically, 'There are no such things as self-organizing systems!' [von Foerster, Heinz (1950). *On Self-Organizing Systems and Their Environments*. Reprinted in *Understanding Understanding*, New York: Springer.] Or, to speak at a conference on ethics and to declare at the beginning of his talk, 'In fact, it is impossible to talk about ethics.' [von Foerster, Heinz (1995). *Ethics and Second-Order Cybernetics*. Stanford Humanities Review. <http://www.stanford.edu/group/SHR/4-2/text/foerster.html>. Retrieved August 25, 2005.] In this way and in many other ways I stand on his shoulders, using this trick, this subterfuge, this moment of surprise, hopefully leading to 'wonder', as Ranulph Glanville has put it in his paper from the proceedings of this same conference.

To start with this trick—which is of course a **second-order** trick—is to jump outside the frame of one's expectations to a new frame. It has the advantage of **reflection**, and reflection is the beginning of **learning**. This is one of Heinz's fabulous insights, elegant and powerful, leading into the nature, and many advantages, of second-order cybernetics.

'Expectations are never met,' as Maturana said only a few days ago at another conference here in Vienna in honor of von Foerster. I hope, however, I will not disappoint you with the path I have chosen today. I believe I can draw references and some specifics about the nature of the work of Heinz von Foerster, Gordon Pask, and Humberto Maturana, and, in this way, weave somewhat of a **mesh** of relationships among these three extraordinary characters, and demonstrate that the thinking of each is inextricably bound to the others.

Outline

Let me give a brief outline of my talk today:

- First, a moment about origins, primarily Heinz's starting point, as context for his **worldview**;
- Then a discussion of my three main points:
 1. Closure
 2. Interaction *with* the environment, and
 3. Interaction *via* the environment.

Note the difference between points two and three. Point two is focused on modeling the relationship *between* cybernetic systems and their environment. Point three is focused on the environment *servicing as a medium* for connecting the system to conversational participants. In the end, these distinctions are solely for the purpose of exposition.

All three of these main points intertwine and so references to each point may occur in any part of the presentation. I will draw references to von Foerster, Pask, and Maturana in regard to each of these three areas. Afterwards, I will offer a birthday present *to* Heinz, as his birthday of 13 November 2003 was only a few days ago. He would have been 92 years old. He said, in a wonderful piece to Niklaus Luhman, 'One should not give answers for the birthday [sic], one should give questions.' And so I promise to offer two questions as birthday presents to Heinz von Foerster.

[Fumbling with the overhead projector to operate it.] And now we will see if this trivial machine will behave trivially. [Pause] They never do.

ORIGINS

I want to spend a moment too—and only a moment—describing how I came to the worldview of cybernetics. This helps explain why I tell the story that I tell in this presentation, and it is an explicit acknowledgment of the autobiography implied in everything that I say, a point Robert Martin so eloquently made in his paper for this conference.

Jerome Y. Lettvin

I first came across the concepts of cybernetics at MIT as an undergraduate, spoken by this man, Jerry Lettvin. [Figure I: Jerry Lettvin in the 1960s]



Figure I. Jerry Lettvin in the 1960s

Here is Jerry in front of a drawing of a vision system, I believe in the frog, the subject of his and his collaborators' seminal papers about animal perception. I am not a physiologist, I am not a biologist, yet I was drawn to this character, to his seductive ways of arguing. He didn't speak the word 'cybernetics' that I recall, but everything he said prepared me to hear it and begin to understand it. Of course, I had heard of the field before, and of Wiener's book of the same name. This field had somehow hit the 2nd-half of the century like a knowledge bomb. My older brother had a copy of Wiener's *Cybernetics* on his college bookshelf, its mathematics unread and unreadable except by math experts. It was still 'in the air' in the late 1960s, long after its ascendance and then fall from grace, but I knew nothing of its history at that point. I hardly knew what the word meant. I thought it had something to do with, yes, computers, robots, and a new era of smart machines. But as Wiener's mathematics was inaccessible, cold, and heavily rational, Jerry's stories were immediate, visceral, and evocative of biological systems.

So taken was I by the style of thinking of the man that I wrote simulations and made a computer-generated movie about axon trees in the nervous system, working with Jerry and one of his research associates, Steve Raymond. That was one way to participate in this way of thinking, by programming their model of nerve spike behavior, not in the neurons themselves (where other researchers had been focused for decades) but in the axons, the transmission cables that wire the neurons together. Everyone else assumed flawless transmission, while Jerry was looking for transformation or computation or filtering performed by the axon as it split and split again, distributing the encoded trains of electrical impulses out across the brain. He and Steve proposed that something interesting happened at every split in the tree. Each axon fiber emerging from a neuron splits into tens-of-thousands of branches, a difficult system to study as pure mathematics, especially without computers. So I wrote a simulation that allowed variation of pulse trains put into the trunk of the tree, as if input from a neuron firing. The graphics screen displayed the individual pulse trains that came out of the leaves, many arborizations later. Watching Jerry's delicate fingers dance on the input switches to the computer, and his delight in its output behaviors, gave me an early taste of the wonders that could lie in scientific enquiries, rigorous and careful, yet intuitive and emotional too.

Humberto R. Maturana

Everyone in Lettvin's classes read one of the seminal papers he wrote with Humberto Maturana, 'What the Frog's Eye Tells the Frog's Brain'. [Lettvin, Jerome Y. and Maturana, Humberto (1959). *What the Frog's Eye Tells the Frog's Brain*. Reprinted in McCulloch, Warren (1965). *Embodiments of Mind*, Cambridge, MA: MIT Press. *Note: Although the paper includes McCulloch and Pitts as authors, this was due to an agreement among the researchers that all papers of this era would include all as authors; however, it is widely acknowledged that the experiments were performed by, and the paper was written by, Lettvin and Maturana.*] This paper casts the frog as a goal-directed system whose nervous system, down to the cells in the retina, were in service of it getting what it wants, for its survival. [Figure II: Humberto Maturana in the 1960s]




Figure II: Humberto Maturana in the 1960s

I read other papers by Maturana along the way, stunned by the concept of 'autopoiesis', as was Stafford Beer when writing his poetical introduction to its early publication. It was much later at conferences that I finally met Humberto and was able to ask him about those early days. Standing casually in a hallway near the registration table for an American Society for Cybernetics conference one evening, he explained how it was because Jerry was out of town, and that Humberto was left to his own inexperience in configuring experiments, that the results retold in the 'Frog's Eye' paper first came to be. Humberto said that Jerry had been struggling for some time, but somehow Humberto got it to work. Jerry later confirmed this story to me.

Andrew Gordon Speedie-Pask

My exposure to these ideas, still without the term ‘cybernetics’, prepared me for what was to come in a way I couldn’t predict. Subsequent to knowing Jerry and reading Humberto I worked in a research lab that Negroponte had started, Nicholas Negroponte, now very famous as the creator of the Media Lab at MIT and a widely read author on technology. One day, a consultant to Nicholas walked into the room, and it was Gordon Pask. [Figure III: Gordon Pask in the 1970s]



Figure III: Gordon Pask in the 1970s

At this moment my world changed. My world changed because here was an individual, a great character with considerable charisma—and craziness too, one might say—whose way of looking at the world was totally different than the **worldview** that was so insistent and insisted-on at MIT. I’d been to classes with Minsky and Papert in artificial intelligence (AI) and had absorbed the literature, and, for me, there was something fundamentally unsatisfying about it. I didn’t know what that was until I came across Pask’s work, and his field called cybernetics. So it was this journey of ideas, through the work of Lettvin and Maturana and then Pask, that led to Heinz.

Heinz von Foerster



Figure IV: Heinz von Foerster at 1962 Aspen Design Conference

[Figure IV: Heinz von Foerster at 1962 Aspen Design Conference] I met Heinz in the early 1980s, when he gave the keynote for a conference on Educational Technology at Concordia University in Montreal. His opening line was, ‘Ladies and Gentleman, I am a Paskian!’ Heinz’s gesture was unabashed support of Pask and his work. Heinz said this, not because it was simply true (Heinz was forever a fan of Gordon, and called him ‘the cybernetician’s cybernetician’) but because Gordon was hoping to obtain a position at Concordia. He eventually did.

Who was von Foerster? You know a lot of this history; I’m only going to retell the pieces that I think are relevant to my focus today.

He was a **physicist**. He used mathematics to **model** the world. He was deeply concerned with quantification. He was deeply concerned with understanding from first principles, from axioms, if you will, how a piece of the world might work. His earlier papers start with a few axioms, simple and clear, and lead to surprising, wonderful, and profound places.

You may know the extraordinary book, *Understanding Understanding* [von Foerster, Heinz (2002). *Understanding Understanding*. New York: Springer.], now available at amazon.com and whose existence is a tribute to many people, especially the oldest of Heinz and Mai's sons, Tom von Foerster, who made sure it was published. In this collection you will find reprints of Heinz's greatest papers, many long-since out of print. There are many examples of Heinz's thinking, of a form and expression that his wife Mai expressed beautifully when she said, 'Heinz has a mind like a crystal.'

The new beauty of this book is that it has new stories. In the preface, Heinz tells of his experience with his own memory as a child. He claims that he was a terrible memorizer and this caused him to write things down when the knowledge to be retained was facts, history, numbers, timing, lineage, chronology. He wrote these things down on a timeline, and he noticed that recent events were very dense on the timeline—many happening in a short period—and events in the distant past were very sparse. So he got the idea of the logarithm of time and, on plotting the events on a log scale, he saw that events were more uniformly distributed. He saw how order could come from models that he applied to his experience.

Later as a scientist he developed a specific and quantitative view of memory—how the macromolecular molecules in the brain, nervous system, and endocrine system behave, and how they create **remembering**. He published this in a paper in German and imagined that no one would read it [von Foerster, Heinz (1948). *Das Gedächtnis: Eine quantenmechanische Untersuchung*, Vienna: F. Deuticke.].

At some point he came upon the 'Ebbinghaus Forgetting Curve'. Ebbinghaus was a character of the 1880s who did experiments with graduate students—we all do experiments with graduate students, that's one of the clichés of talks like this. He derived a curve to show the rate at which, over time, his experimental subjects lost their memories for what he showed them. He plotted a curve of the percentage of recall that would occur after 20 minutes, 1 hour, 8 hours, multiple days, etc. Heinz wanted to compare Ebbinghaus' results to his own, so he took data from Ebbinghaus and tried to match it to his own curve, but it didn't match. He went back and looked at the experiments that Ebbinghaus actually did, and, yes, they did involve initially presenting items to graduate students so that they could remember them. Yet, in order to find out the rate at which the subjects were forgetting these items, Ebbinghaus had to *test* the students at periodic intervals. What constitutes a test for forgetting? The test itself constitutes asking the students to *remember* what had been memorized. But, that act of remembering changes the degree of forgetting. So, Heinz concluded that Ebbinghaus' curve should have been called the 'Forgetting *And* Remembering Curve', because Ebbinghaus was measuring both forgetting and remembering. And when Heinz included that fact when comparing his work to Ebbinghaus, his theory fit Ebbinghaus' data.

I elaborate this story about Heinz's memory because it was critical, when Heinz came to the States and met Warren McCulloch, that *Warren knew Heinz's paper*. And Warren liked the paper by Heinz because it was the first time, Warren claimed, that a theoretical approach matched empirical data.

Experimentalists

A further reason why all this is relevant today is because all these guys—von Foerster, Pask, and Maturana—spent an enormous amount of time in their early careers focused on **experiments**. They were not in their back rooms or front rooms or coffee shops just scribbling and theorizing, they were studying experimental data and experimenting themselves, creating interactions between organisms (or mechanisms) and environments, and extrapolating from there. Experimental work gave them the advantage that they were *actually looking* at the organism's behavior in an environment. This allowed them to avoid predispositions as to what was happening 'inside the nervous system' or what it means 'to learn' or whatever. As a consequence they were able to break out of what would otherwise be the confining and popular **paradigm** of the day.

Heinz used to love to say that once you know that it's a paradigm, it's not a paradigm anymore—because a paradigm is the skeleton *underlying* the way you are thinking which prevents you from thinking in new ways. As a consequence of seeing that it is a paradigm—the second-order step of being outside and reflecting on your model of your experience—you learn that you were restricted by what you *were* seeing. Now you are no longer so restricted, and new theories are accessible to you. So the next time somebody says, 'My paradigm is...' you can reply, 'No, not a paradigm for you anymore...' and you participate with Heinz in this wonderful, magic, mental flip.




Figure V: Warren McCulloch in the 1960s

[Figure V: Warren McCulloch in the 1960s] Warren McCulloch, of course, a physiologist to his core, was doing experiments of all kinds (though let's not forget his study of philosophy at Yale and his wide-ranging capacities). [McCulloch, Warren (1974). Recollections on the Many Sources of Cybernetics. In *ASC Forum*, Volume VI, Number 2—Summer 1974.]

[There is a sudden deep hum in the lecture room.] Is that a sign? Late one night in 1980s I walked across the fields of the McCulloch farm in Connecticut with Gordon Pask, and suddenly a rainstorm arose that was so portentous and surprising. Gordon could hardly keep walking, he was blown about so vehemently, and his cape was flying over his head and almost lifting him off the ground. Gordon told me that he thought this was Warren coming along to meet us. And so, I will say the same thing about this sudden humming sound. [Audience murmurings]

CLOSURE

Here is the start of my first major point, about closure.

Warren McCulloch and Heinz, you can imagine the conversations they would have, given their common interest in nervous system function and matching theory to experiment. One emphasis that Heinz taught me was the distinction between ‘ontology’—the study of the individual—and ‘ontogenesis’, how the individual comes into being. He pointed this out in the question-and-answer session after I offered a taxonomy of some sort in one of my early lectures, and, of course, he was quite correct that I had wanted the latter while I used the former. But never again.

Nervous System Loops




Figure VI: Evolution of the nervous system, from von Foerster

[Figure VI: Evolution of the nervous system, from von Foerster] Along those lines, Heinz reminds us in his papers [von Foerster Heinz (1973). *On Constructing A Reality*. Reprinted in *Understanding Understanding*. New York: Springer.] that the earliest nervous systems had an environmental sensor that was combined with an effector, the part of the organism that enables it to respond with an action in the environment [left section of previous figure]. This is the simplest form of nervous system: sensor to effector, perceiving-to-acting as a single agent.

Then, over time somehow, between the sensor cell and effector cell, new structure evolved. First axons [middle section of previous figure] and then neurons formed in between [right section]. After perceiving, something is being ‘processed’, we might say—while emphasizing that any such word is ‘paradigm-laden’. Any action is taken as a result of the perception *and* the computation done by the intermediate layer. Eventually and in stages, a **nervous system** emerges such that these intermediate cells are doing something.

As we’ve been reminded at this conference, there is a notion called ‘the principle of undifferentiated encoding’. The electrical activity of the nervous system comprises electrical spikes or pulses of activity that traverse the nervous system, coming out of neurons and traveling down axons and impinging on more neurons via synapses, and so on and so forth. Here is a question: What are they saying? What are these pulse trains carrying by way of messages? (Forgive the metaphor of ‘saying’—not to mention the metaphor of ‘messages’—but it’s useful to make the point of undifferentiated encoding.) Well, of course, they’re not saying **what’s** going on, because they are just pulses—that’s the ‘undifferentiated’ part. There might be more pulses or fewer pulses per second, and that difference carries something significant because there’s no other apparent difference in the pulse trains. Furthermore, they are not saying **why** things are happening because they don’t have access to any kind of ‘why’ model; they are only pulses traveling locally here or there. In principle, therefore, all they can say is, ‘**how much**’: there’s this much stimulus, there’s that much stimulus, that’s it. This is the principle of undifferentiated encoding.

Now think for a moment about the nervous system as a whole. The data that Heinz quotes in his paper ‘On Constructing a Reality’—a bombshell of a title if I ever heard one—says that there are about 1×10^{11} sensor cells in the human, if you count the retina and the skin and the hearing system, et cetera, et cetera. But there are 1×10^{16} internal neuronal cells that are not ‘listening’ or connected to what we like to think of as the ‘outside world’, but rather are connected to other things inside, connected to other neurons and *not* to so-called sensory cells. So there are 100,000 times more processing cells in the human body that are ‘talking to each other’, than are ‘talking to the outside’. Now there are many unfortunate simplifications in this way of phrasing the issues, but I think for our purposes here that it’s ok.

As Heinz puts it, ‘We are one hundred thousand times more receptive to changes in our internal than in our external environment.’ [von Foerster (1973). *On Constructing a Reality op. cit.*]

From the perspective of the nervous system, it’s closer to a monologue than a dialogue with the outside world. We can now see the importance of understanding how the nervous system is a **loop**. [Figure VII: From McCulloch’s paper *Finality and Form*]

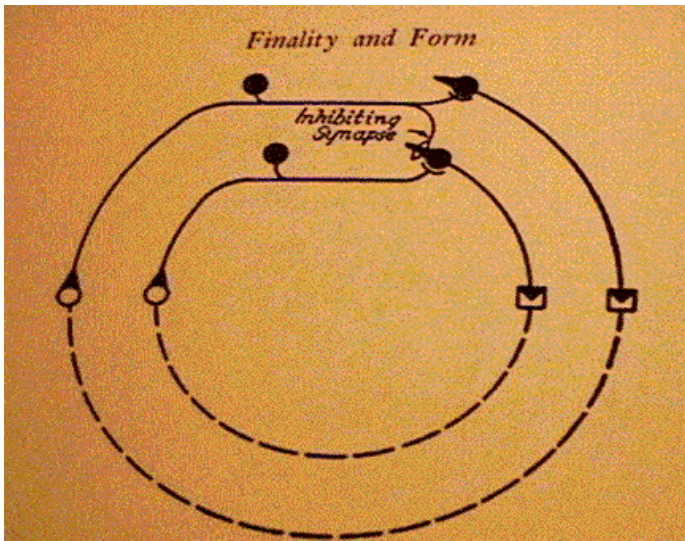


Figure VII: From McCulloch’s paper *Finality and Form*

This graphic is from McCulloch’s paper ‘Finality and Form’ [McCulloch, Warren (1965). *Finality and Form*. Reprinted in *Embodiments of Mind*, Cambridge, MA: MIT Press]. Thinking of the nervous system as a loop leads easily to the concept of **closure** of the nervous system: that everything in the nervous system is nervous system behavior and is not ‘outside-world stuff’. (Those familiar with Maturana’s work with Varela will recognize this concept.) And it should be easy to imagine how one would see the nervous system as engaged in constructing its own reality. Sure, there is a sensory connection *to* the world, but the nervous system is more concerned with what it thinks *about* the world, as implied by that factor of 100,000.

This is not an oversimplification. The notion of constructing one’s own reality **makes sense**, it fits. It **fits the data**. I think one of Heinz’s great contributions is his statement, ‘The nervous system organizes the world to compute a stable state.’ The organism needs stability to function, to maintain its viability, to survive. If the world is not predictable, survival is more difficult—if not impossible. And that stability can be reflected in terms of ‘object constancy’, as with Piaget, or the concept of ‘concepts’, as carefully defined by Pask. We will return to this a little later.

Biological Computer Laboratory

In 1957 at the University of Illinois at the campus of Urbana/Champaign, he started the organization called the BCL, the Biological Computer Laboratory, and a rambunctious place it was.




Figure VIII: BCL envelope addressed by von Foerster

The BCL closed when Heinz retired from Illinois in 1975, which says something about his centrality to its existence. When he wrote to me long after, in the 1980s, he was still using the lab’s envelopes. [Figure VIII: BCL envelope addressed by von Foerster] I love the

juxtaposition of these giant letters of ‘BCL’ and the postage stamp with Hemingway’s image [audience reaction], it all makes sense to me. I know you are familiar with the BCL as an organization and you’ve heard a lot at this conference of the publication, *Cybernetics of Cybernetics*. [Figure IX: Circular cover from original *Cybernetics of Cybernetics*] It is a great compendium of papers and graphics as well as second-order tricks, such its *Meta Book* and *Para Book*—each an inserted booklet, not bound into the main spine of the volume, and whose content exists at different conceptual levels from the main contents of the papers and graphics.



Figure IX: Circular cover from original *Cybernetics of Cybernetics*

In the *Meta Book*, Heinz, in his own inimitable way, published a quantitative analysis of the contributions to *Cybernetics of Cybernetics*. He extracted certain terms relevant to cybernetics (such as ‘information’, ‘systems’, ‘order’, and so on), he counted which authors used which terms, and then he looked at the overlap of terms among authors. He computed what he called a ‘degree of bias’ toward certain terms. He noticed that 10 authors out of 40 referenced the term ‘information.’ Only ten. Nine referenced ‘systems’. Seven referenced ‘order’. He then computes a ‘redundancy figure’ for the papers collected, building on his use of measures of redundancy in characterizing so-called ‘self-organizing systems’. The redundancy figure across the authors in *Cybernetics of Cybernetics* is not a very high number, indicating that there isn’t a lot of overlap or common area of discussion among the papers. Heinz’s conclusion: fundamentally, no one is really anxious to use another’s language. Individuals are ‘in their own world’ more than not. Perhaps this is a manifestation of our nervous systems being concerned with their own viewpoints, and that we have to *want* to share with others for that to happen to any serious degree.

Meta-analysis via THOUGHTSHUFFLER

shuffle VON FOERSTER			
191 texts	CA 04044 04 Report No. 3: Mitotic Indices of Dividing and Differentiating Cells, Heinz von Foerster, February 1961; Report No. 4: Label Incorporation and Label Appearance in White Blood Cells, Alfred Inselberg et. al., December 31, 1963	NSF 17414 TR No. 1 and 2: Linear Property Filters, Alfred Inselberg and Heinz von Foerster, July 1, 1962;	5. Von Foerster, H.: "Basic Concepts of Homeostasis," in Homeostatic Mechanisms, Brookhaven Symposia In Biology, No. 10, U.S.A. Brookhaven Laboratory; Upton, New York, 216 242 (1957).
163 terms			
Von Foerster - 53/191	Von Foerster - 53/191	Von Foerster - 53/191	Von Foerster - 53/191
system			
editor			
Technical Report		Technical Report	
York			York
Ashby			
Computer			
cybernetics			
information			
BCL			
report	report		
Inselberg	Inselberg	Inselberg	
Analysis			

Figure X: THOUGHTSHUFFLER Frame 1

What I want to show you here is a system of my own, called THOUGHTSHUFFLER. [Figure X: THOUGHTSHUFFLER Frame 1] This is a successor to a system called THOUGHTSTICKER which many of you know, originated by Pask and re-implemented as part of the consultancy I managed in the 1980s. [Pangaro, Paul (2001). THOUGHTSTICKER: An Idiosyncratic History of Conversation Theory in Software, and its Progenitor, Gordon Pask. In *Kybernetes*, Volume 30, Number 5/6.] Some of the mechanisms here are the

same, but the user experience is quite different. Think of THOUGHTSHUFFLER as ‘a spreadsheet for ideas’. The spreadsheet metaphor provides the simplicity of rows and columns. I entered the bibliography of the BCL into the system in preparation for this talk. You see each reference as text across the top, and they extend for many, many columns off to the right. There are as many columns as there are entries in the bibliography, 191 in all. Down the left-hand side you can see a few terms that the system has automatically extracted from the bibliography, but I’ll come back to that in a moment after explaining what the you can do with it.

With THOUGHSHUFFLER you can click on any term on the left to select it, and the system *shuffles* the content, that is, it takes all the paragraphs that contain the term that you’ve selected and shuffles them to the left, up against the first column with the list of terms. You can see in this slide [THOUGHTSHUFFLER Frame 1] that ‘von Foerster’ is selected, and the system shuffles to the left the 53 texts out of the total of 191 that contain his name. Now this is a kind of ‘search function’, because the system displays the results of a search for a specific term. But it does more than what the usual search engines do.

191 texts	152. Von Foerster, H., A. Inselberg & P. Weston: "Memory and Inductive Inference" in Bionics Symposium 1966: Cybernetic Problems in Bionics, H. Destreicher & D. Moore (eds.), Gordon and Breach Science Publishers: New York, 31 68 (1968).	231. Weston, P.E. & H. Von Foerster: "Artificial Intelligence and Machines That Understand", Annual Review of Physical Chemistry, Vol. 24, H. Eyring, C.J. Christensen & H.S. Johnston (eds.), Annual Review Inc.; Palo Alto, 353 378 (1973).	C M D F R I A C D
163 terms			
Von Foerster - 53/191	Von Foerster - 53/191	Von Foerster - 53/191	V
editor	editor	editor	
York	York		
cybernetics	cybernetics		
Inselberg	Inselberg		Ii
science	science		
Weston - 9/191	Weston - 9/191	Weston - 9/191	
memory	memory		

Figure XI: THOUGHTSHUFFLER Frame 2

Now imagine that I click on ‘Weston’. In the next frame [Figure XI: THOUGHTSTICKER Frame 2] both ‘Von Foerster’ and ‘Weston’ are selected. THOUGHTSHUFFLER will take the bibliographic entries that contain both of these names, and shuffle them to the left. Here we see the only two entries with these two authors in common. In this way, the user manages the display of content in the spreadsheet, focusing down and de-focusing on the many terms.

Some of the THOUGHTSHUFFLER’s functions are very simple, and others more subtle. For example, it can facilitate ‘compare and contrast’ across multiple paragraphs, whose similarities and differences are manifest by the selected terms and juxtaposed columns of similar text. [Pangaro, Paul (2001). *The Nexus of Search, Models, and Meaning*. <http://pangaro.com/thoughtshuffler/>. Retrieved August 25, 2005.]

One can use THOUGHTSHUFFLER as a front-end to search engines, to web sites, and to a variety of content sources. I propose it here as an *homage* to Heinz’s interest in extracting whatever can be extracted from the low-level syntactic or structural components of a language system. I mean this not in the lineage of ‘natural language processing’ (involving AI notions that presuppose that the meaning is *in* the text), but rather in the sense that the structure and details of content may be ‘mined’ usefully. (This does not controvert the cybernetic tenet that meaning is created by the hearer and is *not contained* in the utterance itself; rather, everything about THOUGHTSHUFFLER manifests the commitment to believing that meaning is made by the listener, who reproduces the entailment of specific distinctions that cohere to **make sense together**.)

This ‘mining’ approach is especially manifest in automatic processing by THOUGHTSHUFFLER that uses Pask’s entailment notions to self-extract phrases that could be of interest to a user—someone wanting to understand, for example, what did von Foerster say about cybernetics? How often did von Foerster and Weston collaborate and what was it about? In the broadest terms, THOUGHTSHUFFLER aids the user in asking the question, what is the coherence of this textual content as a body of knowledge? I find it fun to watch what it’s doing as it automatically extracts phrases that it computes from Pask’s logic of coherence and distinction. [Referring to the running software, now automatically extracting new terms from the bibliography] Here it’s saying that ‘superposition’ is important in 5 of the 191 references. It would be wonderful to have a conversation with Dr. Paul Weston about how his work in ‘cylinders’ might be applied under this scheme.

Now I shall have to stop THOUGHTSHUFFLER or it will rob all of the processing power of the entire computing device. [Halting the calculation and returning to display of figures from the paper]

Common Concepts

I want now to ask, what was the degree of bias across the work of Heinz von Foerster, Gordon Pask, and Humberto Maturana? How much were they interested in the same concepts? To my knowledge no one has done this analysis.

They certainly used a lot of the same terms. You will find the words ‘order’ and ‘coherence’ across their individual works. ‘Concept’, ‘interaction’ ‘organizational closure’, ‘memory’, ‘structural coupling’, were used often by each and, I believe, borrowed from one to another. You could check the dates of the papers to attempt to find the first published instance and then say, it first appeared here used by Humberto, so Heinz got it from that guy. In my view, although the phrases are the same, the underlying concepts, what those concepts **entail**, are different among them, at least in the sense of the *scale* at which they were applied. Heinz was interested in talking about memory in terms of its convergence and stability, as modeled by mathematics. Gordon was interested in talking about memory at almost a mechanical level. Humberto was interested in the level of the social, the language, the ‘living together’. These are similarities and differences that I want to draw out for you and not elide.

I believe, however, that when they used the word ‘coherence’, they mean the same thing, their interpretations are the same, and, if we had them all in the room at the same time, they would—if they were willing to have such a conversation—agree to very strong overlaps. Before I go into detail about what is common, I want to make another observation.

Heinz, as many of you know, was a very generous character when it came to ideas, and he always referenced others. Heinz makes many, many references to Pask and to Maturana. For his part, Pask very often referenced others. I hesitate to say he was generous. Surely he gave genuine credit to others in certain cases. But I think he liked to play the game ‘my world subsumes your world’, and by referencing others he could incorporate their work as a subset of his own. In this I do not think he was inappropriate, for I think his theory *is* comprehensive and original. Moreover, he referenced others in hopes that readers would see connections that he himself was not the most lucid at making.

Finally I point out that Maturana does not reference others. [Audience murmurs] This may seem a lack of generosity. However, Heinz himself has the answer. Heinz says, in essence, that biology must write its own laws. And for me this absolves any need for Maturana to reference others. [Audience laughter]

INTERACTION WITH THE ENVIRONMENT

I want to give a marker for where I am. I am moving from my first main point, closure, to my second, that of interaction with the environment. This is a somewhat technical part of the talk. We had the demo, meant to be amusing, diverting. Now, I have a bit of deeper content, if you will allow me.

'Looping-through' and Autotopic Mapping

In the 1980s I wrote a monograph called 'New Order From Old' [Pangaro, Paul (1988). *New Order from Old: The Rise of Second-Order Cybernetics and Implications for Machine Intelligence*. <http://pangaro.com/NOFO/>. Retrieved August 25, 2005], and therein described my journey through second-order cybernetics. In one part of the work I was trying to describe the nature of 'looping-through', as I called it—the nature of the recursive activity of organisms and the environment, the loops that the organism completes through the environment rather than those that it completes through its internal structure, as we were talking about earlier. I was struggling with that. I didn't have a better term than 'looping-through', with its awkward hyphen.

When I sent the paper to Heinz for general comments, he sent me back a postcard. Written in his unmistakable hand, the postcard held the words, 'Paul, many wonderful points!' Which led me to ask myself, which were the points that weren't so wonderful? [Audience laughter.] He had such a frankness. I'm sure that if I asked him, he would have told me precisely what he thought. He had a way of being very, very positive, and very constructive and critical all at the same moment. I always appreciated that with him, from him.

Without any specific reference in his writing on one side of the postcard, the other side held the following figure. [Figure XII: Postcard from von Foerster with taxonomy of mappings of organism and environment]




Figure XII: Postcard from von Foerster with taxonomy of mappings of organism and environment

By sending these figures, Heinz showed me that he saw my struggle and that *he* had a very precise model for looping-through, part of a complete taxonomy of mappings between organism and environment. His name for the 'looping-through' mapping was 'autotopic', meaning 'self-mapping', shown in the bottom section of the postcard. The organism on the right of the bottom section is shown having a nervous system whose function is to map its experience with the environment through itself and back to the environment (A to A'). This is in contrast to the mapping in the top section, which is the AI paradigm, the 'input/output' paradigm, where one says that there is an objective something 'out there', and that the nervous system has a model of the outside 'in here', inside the organism, and that those internal and external elements are in 1-to-1 correspondence. (The middle one simply completes the space explored by the taxonomy, showing a third theoretical mapping.) This was a great moment of clarity for me, with Heinz coming forward with this beautiful and clear concept. I can only imagine all the times this must have occurred in daily conversations at the BCL.

Von Foerster's Eigen Functions vs. Strange Attractors

Moving from this modeling and mapping, I want to draw something that I expect all of you have seen. Here is a way of describing the **convergence** of the nervous system on a stable autotopic mapping as a consequence of interaction with the environment. [Draws the next figure, a converging Eigen Function] [Figure XIII: Eigen Function that converges to a stable point]



Figure XIII: Eigen Function that converges to a stable point

Heinz not only connected the ‘forgetting-and-remembering curve’ to a model of learning in the nervous system, but he also made a major contribution by offering a mathematical model to embody the idea that, when computing a stable state, the nervous system converges on stable things called ‘topics’ or ‘objects’. This can be drawn as a convergence in state space, a spiral that moves and converges onto a consistent, coherent, stable point: the blob that I’ve just drawn. This is a ‘stable point’, an object, a persistent concept of ‘chair’ or ‘table’ or ‘circle’ or ‘freedom’. I want to say again that it is a stability that exists as a consequence of interaction of the organism with the environment. I asked Heinz where the idea of modeling concepts and memory as Eigen functions came from, and he said, smiling, ‘Well, I suppose it came from me!’ But I thought, no supposing about it.

You probably recognize this diagram as one that is very popular in fields that arose long after ‘self-organizing systems’ and cybernetics, long after when Heinz was developing the concept. The diagram is used in chaos theory and complexity theory, where this function is called a ‘strange attractor’—‘attractor’ in that the system appears to be attracted to a particular state or convergence; ‘strange’ in that we don’t really know why or where the attraction comes from. This is an unfortunate paraphrase of the concept of ‘system converging on a goal’ where the goal is stability of experience, also known as ‘learning’. And it was always great to hear Heinz say, ‘strange attractttt-orrrr’ and to see by his knowing smile that he was taking credit for having generated the concept long before it became popular. (Heinz also rightly took credit for the massive popularity of G. Spencer Brown’s book *Laws of Form* because of the review he wrote for Stewart Brand’s *Whole Earth Catalog*—a popularity wave that I experienced, finding it everywhere in the middle 1970s. There were so many things he did originate, without credit, though he never expressed anything even so mild as disappointment.)

Pask’s Concept of ‘Concept’

I want to draw a very specific analogy from Eigen functions and Pask’s notion of ‘organizationally closed concept’ and I ask you hold on to your seats, it may be a bumpy ride.

A stable concept represented as a convergence to a small area, as per this diagram, is a simplifying, high-level view, but an idea or a concept is a complicated thing. It’s not a blob that has converged to a stable point in some 2-dimensional state space, that’s merely a metaphorical representation. A deeper model must be available, one that decomposes a concept into components. Pask wrote a great deal about the sub-parts and processes of concept repertoires. [Pask, G. (1981). Organisational closure of potentially conscious systems. In *Autopoiesis*, M. Zelany (Ed.), New York: Elsevier, North Holland, pp. 265-307.] He developed a vocabulary and structural model that included well-specified meanings of terms such as ‘topics’ and ‘entailments’ and ‘coherence’.

The canonical example is the following: if I wanted to explain to you a new idea, a new **topic**, I could do it by calling your attention to other topics that you already know. For example, you know what a circle is. (For the sake of the illustration, I presume that we do share a common meaning for ‘circle’.) And you know what a flat surface is, like this table, or a plane, or something like that. Well, what I’d like to do now is tell you about a new topic that I don’t believe you know, and that new topic is an object, a thing, and what we’re going to do is use this new thing in order to draw a circle on the table. That is the goal, that is the **why** of what I want to tell you. That’s the first part of our exchange about this.

In the second part of the exchange, I'm going to actually tell you **how** this all works together. It looks like this [gesturing with hands to show the arms of a compass that is used to draw a circle], and it has something here that moves [indicating the hinge that sets the diameter of the resulting circle]. I can put one of its points down [gesturing again to place a point of a finger on the table], and make sure it doesn't move around on the table. Now I'm going to swing this other end and make a mark and move it around on the table and leave the figure of a circle [continuing to gesture].

So you watch me do this dance, and you see me mark something on this table, and you recognize that I've drawn a circle on the table.

Let me tell you that this thing I've used is called a 'compass'.

Now, at a high level this is an example of Pask's notion of an 'organizationally closed concept'. There were three stable elements involved—circle, table, compass—and let me call them in the abstract T, P, and Q, following the conceit of Pask's writing. Notice that these things are distinct from one another—a circle is not a table, a table is not a compass—but they exist in a very tight relation that Pask called a 'coherence'. Let me tell you what is required to have a coherence.

I'm going to remind you now that our exchange had two parts. I said, here is **why** I want to put these things together—the relationship among these topics and the roles that each will play. I'm going to use this thing to make one of those things, a circle, on a third thing, a table. But **how** it all goes together is separate, involving the mechanics. Although he later changed his terminology for these two complementary aspects, in order to make them more general, we can usefully call them the 'description-building' and 'prescription-building' aspects, respectively.

The reason why I wanted to make sure that we realize that there is a descriptive part and a prescriptive part, and that all of these things come together in a particular way, is because Pask would then model the coherence as the following set of relationships. [Figure XIV: Pask's notation for participant A understanding T in terms of P and Q]

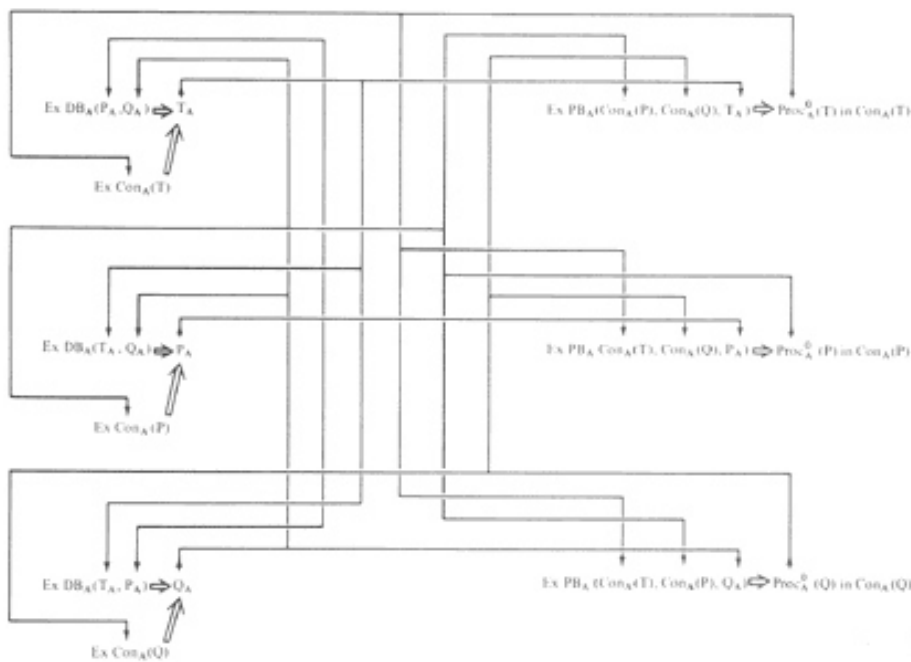


Figure XIV: Pask's notation for participant A understanding T in terms of P and Q

This, Ladies and Gentleman, is the minimum set of processes required to reify a concept, to create a stable point, to express an Eigen function in the conceptual repertoire of a participant called 'A'. It consists of three distinctions or topics, notated here as T, P, and Q. Looking briefly at the top line of the three lines in the figure, you will see on the left side (simplifying the notation for a moment) that $P + Q$ can produce T. This says that combining the descriptive aspects of P and Q leads to the descriptive component of T. This is Pask's notation for the previous notion that 'a circle can be drawn on a table by using a compass'. On the right side of the top is the corresponding notation for the prescriptive aspect, namely, how the elements go together to actually explain a compass; I won't draw those details out any further. Your head may be hurting and I'm not going to go much deeper, but I believe that it was the notions of a circular nervous system that converges on stability that led Pask to seek a model at a lower level, or a deeper model. These two models—Heinz's Eigen functions and Pask's coherence—are isomorphic.

If you care to follow the details of the figure, you will notice that any of the three topics is produced by processes that combine the other two. Thus, no topic is of 'higher order' or 'more abstract' or 'bigger' than any other; all three have equal status. There is no

hierarchy here. I'll leave you to further explore coherence models directly from Pask's writings. Now I'm just going to hint at the deeper direction because anytime you can get a plug in for Pask that makes sense from the context, it's a good thing.

The previous figure is the top half of this next figure [Figure XV: Pask's notation for participant A and B having different understandings of T]; again, for example, the understanding of circle, compass, and plane that I just described, here notated as T, P, and Q, in the repertoire of participant 'A'.

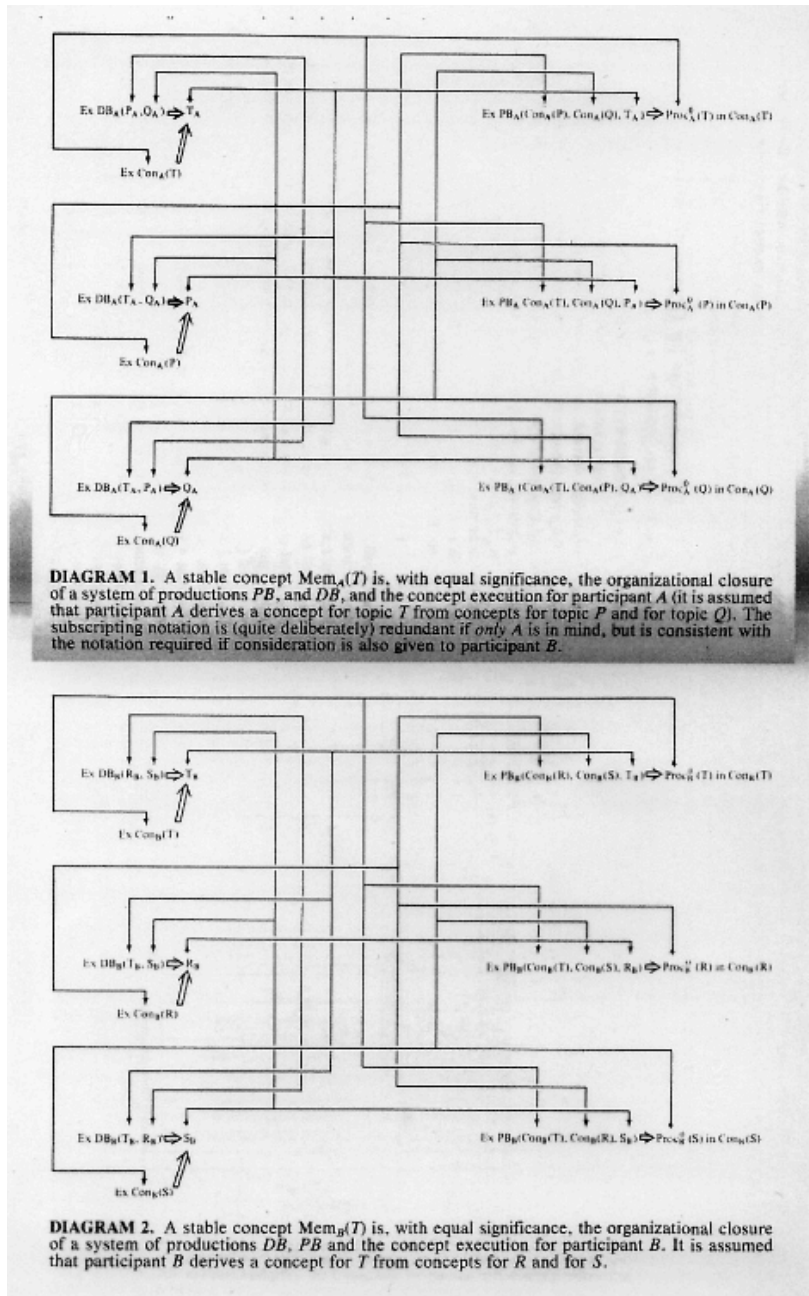


Figure XV: Pask's notation for participant A and B having different understandings of T

But there may be another individual I want to converse with, and not someone who understands circle in terms of compass and plane, rather someone who takes a cylinder and makes a very thin slice perpendicular to the length. This is a different way of entailing the concept of 'circle'. The lower part of the figure represents this second conceptual repertoire, identical in form to the previous one but relating to participant B and where the relation is that of T, R, and S, symbolically representing 'circle', 'cylinder', and 'slice'.




DIAGRAM 3. *L* agreement over common understanding of topic *T*. *A* derives *T* from *P* and *Q*. Participant *B* derives *T* from *R* and *S*. An agreement may be complete or partial depending upon the isomorphic part.

Figure XVI: Pask's notation for A and B coming to agreement over T

Pask then draws out all of the interactions in the next figure [Figure XVI: Pask's notation for participant A and B having different understandings of T]—don't worry about reading the details, just get a sense of the complexity and inter-penetration of all the descriptions and prescriptions. This is the **process of coming to agreement**, which he also calls **synchronization**—a great word, very important to be reminded of in Pask's **worldview**, because he talks about conversation as synchronization of *a priori* asynchronous processes. It is unfortunate that there is not more time today to delve into the poetry and power of that view.

If you do all of that—if you and your partner in conversation transact sufficiently to trigger the synchronizations and productions written in the diagram—then you end with the structure of the next diagram **being shared** by each of the participants [Figure XVII: Pask’s notation for participant A and B having different understandings of T]. It is important to note that not every line of production in the figure requires an individual interaction, utterance, or exchange; only that the conversation carries *enough* triggers to start a process that, *because of the structures and constraints inherent in the specific descriptive and prescriptive aspects of the topics*, results in a stable, common **understanding** of the coherences that result from the combined worldviews.




DIAGRAM 4. Given that *T* is learned by one participant (for consistency, participant A) who derives a concept of topic *T* from concepts for topic *P* and *Q* or *R* and *S*, the stable concept is an organizationally closed system of productions that take place in one participant.

Figure XVII: Pask’s notation for A and B’s shared understanding of T in terms of both R and S, and P and Q

Let me paraphrase this last section about Pask’s model of coming to agreement. Before our conversation, I thought one way (T, P, Q), you thought a different way (T, R, S). We have a series of transactions, we do our little exchanges, we engage in this magic thing called language, and we fill in the gaps ourselves and we **make sense of it** individually, having been synchronized by the common elements. Conversation works because of the deep constraints of the elements in play at any given time, and the small number of ways that the individual components may combine. (We are all aware of the misunderstandings that can arise, but it is a wonder that these are so few; conversely, it is a wonder that so much is comprehended and shared so easily.) In the end, we both possess the process-interaction structures modeled in the previous figure that we can re-activate at any time. This is called **knowing**. We both can

construct T from P and Q, or construct T from R and S. And, quite importantly, the concepts each of us possess are *consistent* or *compatible* or *synchronized*.

I may say to you later, ‘Can you recall for me please, what do you think of when I ask you about a compass?’ That action that we call ‘recalling’ is not, as Heinz showed in another of his wonderful papers [von Foerster, Heinz (1965). *Memory without Record*. In *The Anatomy of Memory*, D. P. Kimble (Hg.), Science and Behaviour Books, Palo Alto, pp. 388–433.], the **retrieval** out of some bookshelf of this little piece of knowledge that I pull out and then I narrate to you. It is the **recomputation** of how you know what you know, how you **make sense of it**. It is that last diagram from Pask that shows the processes that are performed by the nervous system in order to do that recomputation, to ‘remember’ what we know. Pask also called this process, the **production and reproduction of concepts**. Von Foerster shows that remembering must be recomputation, arguing from axioms; Pask develops a detailed model of how that takes place.

A final comment about Pask’s models of concept formation: Each of these diagrams is a spectacular example of **closure**, the first main point of my talk. The structures and processes fold back on themselves, representing the recursive and self-reinforcing nature of nervous system processes. I don’t know anyone who has gone to this level of specificity and clarity about the nature of cognitive understanding and coming to agreement. We don’t know how the nervous system embodies these structures, but I’m highly confident that Pask’s models move our knowledge forward in a great leap. I believe that these Paskian models provide specific directions for neural experiments, and can influence the most modern efforts to make intelligent machines, as for example by Jeff Hawkins’s company, Numenta.

Maturana’s Niche Model




Figure XVIII: Maturana’s drawing of the niches of A and B, including environment labeled C

Let me cite an example from Maturana at a different layer of abstraction, a different kind of abstraction. The example that I showed you from Pask asserts that when we come together in conversation, we each have our own individual experiences—our personal niche that is the basis of our understanding the world. As a consequence of that coming together, we emerge in a different condition, if you see what I mean; we now share something we didn’t share before in the form of an understanding. Similarly, here in Maturana’s diagram [Figure XVIII: Maturana’s drawing of the niches of A and B], A and B have different niches, different prior experiences, but as a consequence of their interactions, they **now share** the common knowledge of **what they mean** by circle, compass, plane, or circle, cylinder, slice. I believe that this is precisely the same concept, shared by Maturana and Pask.

One could say a lot more about this. For example, how do distinctions such as circle and table arise in the first place? Why does there arise the ability to see the difference between this and this and this, or this and that, or me and you. That’s another story for which Pask also has explicit models, but it is beyond today’s scope.

I like very much the quote from Maturana, ‘In recursion, something new arises. Only in recursion can we have an intelligence’—because if you’re not doing it again and again on top of what you had before, and changing as a result, you’re not going anywhere in the state space of learning. You’re not converging on anything, you’re not computing a stable state of the nervous system. You’re dancing along but not evolving.

Pask’s Musicolour Machine

Again I allude to ‘autotopic mapping’ from a few minutes ago, the postcard from Heinz with the mapping from environment through organism through environment. Pask, ever the experimentalist, built machines and created interactions between machines and humans, and even between machines and machines. There’s a wonderful Pask device that I’ll spend a minute describing, trying to recreate for you, called ‘Musicolour.’




Figure XIX: Architecture of Pask's Musicolour machine

[Figure XIX: Architecture of Pask's Musicolour machine] Imagine a musical keyboard and a musician playing it. [Referring to details in the figure] That's meant to be a hand with a finger touching the keyboard, that's not very good, is it? But I think you get the idea.

The keyboard sends a signal representing the music into a box that has a comparator function that compares two inputs and generates an output. One of the inputs is the music and the other input is a memory of what came before. The comparator sends a signal to a response generator that drives a bank of lights. The musician sees the lights, projected onto large hanging sheets of fabric, and this closes the loop from music to lights to musician to more music. That's the organization of the Musicolour system.

This was not what we may remember from the 1960s, those of us who do remember the 60s, where projected lightshows sprang up in dance clubs and home living rooms, driven by recorded music. In those instances, the function of the system was to map sound to light. Bass sounds would show up as red or whatever, and treble as yellow, and so on. There was no feedback from the lights or any human. In sum, the mapping of aural to visual was constant. Because nothing novel happened, it was quickly uninteresting, and the fad quickly died out.

Musicolour, on the other hand, was a cybernetic system: it had a goal, and its primary goal was not to be *bored*—this was Pask's insight, to construct such a system and to couple it with a human performer. If I'm playing my instrument in a particular frequency range, and Musicolour is listening in that range, then it shows me these nice colors projected and dancing in synchronization to my playing. I'm having a good time for the moment because I can make the lights dance. But if I continue to play in that same range for more than a short span of time, the system gets bored with my playing and slowly shifts where it's 'listening'. Let's say it shifts its listening to higher frequencies. I'm not playing in those frequencies, so the system seems to be less and less responsive and doesn't show me any pretty lights in synchronization to my playing. Not much is happening, I start to wonder why I'm not getting any response, and I begin to explore and play somewhere else. If I move lower in the scale, still nothing happens; if I play higher in the scale, the system responds again and I see the lights and I respond in turn. And so on, in this loop—as the system moves in its attention, I move in my attention, and vice-versa. Together, the outcome is an ever-changing dance. Both participants, musician and Musicolour, have goals and they interact with each other. That loop is the autotopic loop where the organism—the musician, in this case—senses the environment, maps through itself and takes certain actions that map back out to the environment. (You could equally well speak from the perspective of the device, where Musicolour senses the environment of the music, and takes action in the form of the light show that maps back out to the environment.)

Pask and his colleagues (including Elizabeth Poole, whom he later married) traveled around England in two small vans, installing Musicolour in music halls and drawing the public in to these shows. [Pask, Gordon (1970). *A Comment, a Case History and a Plan*. In Reichardt (Ed.), *Cybernetic Art and Ideas*, London: Studio Vista pp76-99.]

This was I believe the very first experiment of Pask's, done in the 1950s, a physical manifestation of the general nature of experience, if I may put it that broadly. This corresponds to Maturana's niche diagram from a moment ago, with interactions now construed to be between human and machine sharing the same niche.

In these examples, I'm focused on interaction **with** the environment, to explain how these three different cyberneticians would describe it. That, Ladies and Gentleman, is my number two main point.

INTERACTION VIA THE ENVIRONMENT

My third major point, the nature of interaction *via* the environment, begins with a brief story.

The Constructivist and the Physicist

I worked with a physicist in the 1980s building a large-scale version of THOUGHTSTICKER software based on Pask's ideas [Pangaro, Paul (2001). THOUGHTSTICKER, *op.cit.*] My physicist colleague, a brilliant guy by the name of Jeffrey Fancher Nicoll, couldn't get the **constructivist position**—it didn't make sense to him that every individual constructs a world and invents its attributes, rather than discovers the objective properties that the world *a priori* possesses. He was upset about it and we argued about it frequently. Finally he came to me in triumph one day because he thought he had an example that would win me over to his position. He said, 'Paul I'm going to walk across the room and when I get to that chair, I'm going to trip over it. And when you walk across the room, you're to trip in the same spot—and that means that the chair has to be there!' And I thought, oh no, he got me! Maybe this proves the world really does exist and its properties can be known, what am I going to do? And I realized, wait, Heinz was trained as a physicist, maybe he can answer this physicist's question. So I called Heinz and I told him the story and I asked, what are we going to? And he said, in that distinctive thrust of voice, persuasive because of its energy alone, 'Paul, when you walk across the room, you both feel pain! From there, everything is inference!'

That's all I needed. What can you say to that? [Audience laughter.] Everything else *is* inference, the only thing you have is the pain, the perception, which is the extent to which you have access to whatever you might claim is 'out there.' Constructivistically speaking, the conversation that emerges from this point of experience is *all* inference.

Pask's 'Theory of Everything'

The equivalent of that story is this next diagram.




Figure XX: Pask's Conversational Architecture

[Figure XX: Pask's Conversational Architecture] The genesis of this diagram meshes with my story here today. Gordon was asked by Negroponte to write a paper that would address the construction of an 'architecture machine', a computer environment that would interact with architects in the design process. The result would be like architects working together rather than architects working with machines. Pask wrote up a theory that had been percolating at his lab for a few years, called Conversation Theory. When I first came across it, I called 'a theory of everything'. I hope to give you a sense of why.

Let's have a look at the diagram, though I won't explain all the details here. The broad horizontal and vertical lines are in the original drawing, but were re-drawn in very thick marker by Negroponte before he published them in his *Soft Architecture Machines* book—it was a brilliant stroke to use Gordon's original drawings, rather than have them re-rendered in some sterile style, and this was also consistent with the meaning of his book.

The vertical axis separates 'alpha' and 'beta', two humans in a conversation. These correspond to the 'A' and 'B' subscripts that appear on each side of the vertical, referencing individual participants just as earlier figures have done. Notice the horizontal loops going across the upper and lower parts. We can think of the stuff that traverses these horizontal loops as triggers of attention to, say, look at this circle, look at this compass, see how these interact here, etc. These are triggers into the structural, closed nervous systems of the repertoire of each participant. Earlier I was describing all the inter-process synchronization diagrams from Pask. One can think of those as the details of the effects of the triggers that loop horizontally in this diagram, bridging the interface between human participants A and B. As previously noted, the triggers are carried by language and hence Pask calls these horizontal loops 'I/you-referenced interactions', because the interaction involves treating the other participant as a peer ('you') in relation to the 'inter-actor' ('I').

Although these triggers occur at multiple levels in the diagram, separated by the strong horizontal line, there is a difference in the relative meanings of the upper and lower levels of loops. There are many ways to characterize this difference, but for our purposes here I will say that the upper loop occurs in language or discourse at the level of **why**. Take the earlier example where I said that I will be using this thing (the 'compass', as you learned later) to draw a circle on the table. This is the 'why' of the conversation, and **what may be known** about the domain. Earlier I called this the 'description-building' part of the interaction.

Correspondingly, the lower-level loop is an interaction in language on the level of **how**, or using a different wording, **what may be done**. This is level at which I describe how I put this point on the compass down, move it around, make the mark on the flat surface, and so on. This is the 'prescription-building' part of the interaction previously described. Note that these various ways of describing the difference of the upper and lower loop are somewhat informal and meant to be evocative. Although found in Pask, this interpretation is not so formal as the notation of the original paper from which the diagram is taken.

Let's look now at the vertical interactions, going from the upper half (above the horizontal line) down to the lower half, and then looping back up again. These interactions are **control** interactions, called '**it-referenced interactions**' because the inter-actor is acting to control the other, treating the receiving end of the interaction like an 'it' or an object and not a peer. If inside myself I say, 'Paul, walk to the overhead projector', that's like me up here [pointing to upper-left quadrant] saying to that down there [lower-left quadrant], 'Walk to the overhead!'—but I'm not actually *saying* it, I'm literally controlling it; my legs don't have any choice.

You can model these vertical interactions in many, many more layers than I just called out. You can look at the layers of the zips and zaps going down my legs to control my movement—my brain controls my spinal cord, my spinal cord controls my leg muscles, etc.—or just think about it in the previous abstract form, at a level of description for purposes of exposition.

But here's a different approach to getting what I want: instead of controlling myself to take a specific action, I can come over to Robert, for example, and attempt to control him. [Moves to Robert Martin and gestures as if to take his arm]. I can take his arm and move it down. That seems like a nasty thing to do, and in many cases it is. But there are situations when he's fine with that. He's fine with something coming from me here [upper-left quadrant of the figure] as a control function operating down on him there [lower-right quadrant] so long as he says to himself [tracing a line from lower-right to upper-right], 'Oh, let me think about what just happened: Paul's doing that because he's trying to demonstrate a point, that's fine with me'. It's acceptable because we share a common goal at this level [upper-left and upper-right], the level of **why** we're doing it. We're doing it because we both agree that it is reasonable to do such a demonstration.

Another example that explicates the figure is **dancing**, perhaps the best example of all and one used by Gordon himself. I'm dancing with my partner and I shove her this particular way as we dance. She doesn't think, 'How dare you?!' She thinks, 'Well, why is he doing that? Ah! Because we are dancing.' And, in fact, when she gets shoved by me and realizes that up here my goal is to dance [upper-left quadrant] and she also has that goal [upper-right quadrant], she shoves me back! Fine with me! I'm willing to be shoved, because I'm willing to engage in this for the common goal. And in the end this criss-crossing [gesturing in a figure-8 across the two participants horizontally and across the two levels vertically], this back and forth, this intertwining is two systems, an A and a B, an alpha and a beta, agreeing to have a common goal and **to behave like a single system**. In that merging is agreement, cooperation, and a willingness to give myself up my individuality.

This one model encompasses objective interactions (it-referenced) and subjective interactions (I/you-referenced) in a single frame.

Maturana's 'Do You See The Flower?'

I can spend hours talking about this model, which is for me one of the most profound cybernetic models ever drawn. This is analogous to Heinz's story where we trip on the same chair: we do a dance in our perceptions, we have a common experience that we somehow wish to explain, and we use language to do so. This is summarized in Humberto Maturana's drawing [Figure XXI: Maturana's drawing of shared experience in terms of inter-twining nervous systems of A and B].

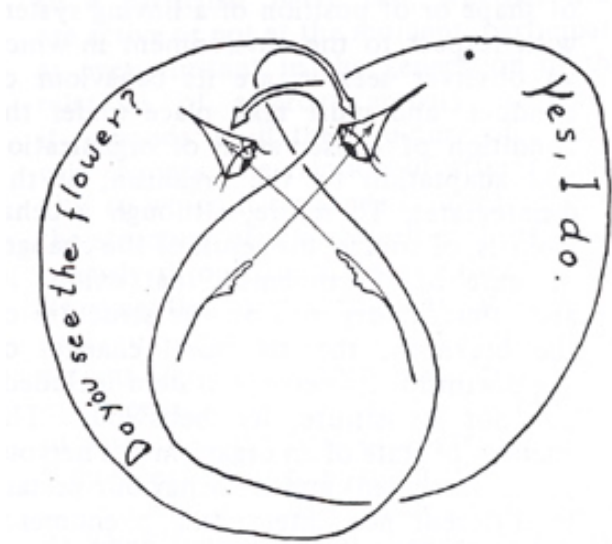


Figure XXI: Maturana's drawing of shared experience in terms of inter-twining nervous systems of A and B

We see two nervous systems, which of course don't have direct access to each other—my nervous system is closed, your nervous system is closed. But there are triggers back and forth that occur through the sensory networks, and (skipping a lot of intermediate steps and metaphorically speaking), it is as if our nervous systems intertwine and we agree, 'Ah, what a beautiful flower!'

Maturana's drawing is about 'recognizing the legitimacy of the other', as he phrases it. So too is the previous figure from Pask, where we impose 'control interactions' on others and allow control interactions to be imposed on us only when their goals are compatible with ours and when, in reciprocal respect, we agree to interact, to converse, to move from asynchronous to synchronous interactions. After we each trip on what might be called a chair and both feel pain, the manner in which we engage is an ethical question that is richly, poetically explored by Heinz in his 'Ethics and Second-Order Cybernetics', my desert-island paper [von Foerster, Heinz (1995). *Ethics and Second-Order Cybernetics, op. cit.*].

This was my main point number three, interaction via the environment, drawing examples from each of these three extraordinary characters, pointing out some of the many commonalities in their work.

Let me review my three main points. Von Foerster, Pask, and Maturana share the view of the nervous system as a closed loop. They share the constructivist position that an organism's nervous system builds a repertoire of stable interactions with and through the environment that are compatible and resonant, and these interactions help the system maintain viability. And they share the understanding that these interactions constitute 'living together' in social relations, woven by conversation, whereby common goals are possible, collaboratively derived, and productive.

I am excited to see a surge in interest in cybernetics, both its history and its application, and I hope to see deeper investigations into these rich connections I have sketched for you today.

STILL TO EXPLORE

I have just a couple of comments and then my two questions for Heinz's birthday. I want to acknowledge some points that are not explored in this talk but that are important to my theme.

Pask was obstinate in distinguishing between a learner's 'style' (the internal, unknowable processes by which a learner comes to understand something new) and a learner's 'strategy' (the exteriorized evidence of the path the student took to learn successfully). For

him, it was an ethical requirement not to presume that you can know what another is thinking; we have only the evidence of what is exteriorized.

At a higher level of abstraction and importance, and consistent with Pask, Heinz was original and disciplined in his work on trivial and non-trivial machines, showing that very simple mechanisms could be non-trivial—that is, unpredictable—while being deterministic. This stood for him as an explanation of how a human being, while being deterministic, could be both unpredictable and without the need for any agency other than the mechanism of the body to explain all its wonder.

The scientific community largely misses Pask’s use of the ‘null-point hypothesis’ to structure experiments that extend into subjective realms. His approach leads to the ability to hold objective (vertical) and subjective (horizontal) interactions in the same frame. He can thus universalize *all* kinds of interaction into a single, rigorous model of ‘conversation’, and in this way he has produced ‘a theory of everything’.

Pask’s rigorous view of science requires any worthwhile theory to contain laws of complementarity, duality, and conservation, and here he hooks deeply into the scientific tradition. Maturana speaks eloquently of the scientific method and how it does not require an objective reality to work—an extraordinary view. And I’m reminded of Heinz’s pointing out that ‘science’ comes from the same word as ‘schism’. Science works—and here he means the usual scientific disciplines of physics, biology, sociology and all those other ‘ologies’—because science divides a complex world into smaller and smaller pieces, until finally you are able to say, ‘Aha, A causes B, I’m pretty sure!’ Dilemmas arise when you come across problems that require a holistic science and you want to say, ‘A causes B causes C causes A’ and you get tripped up; hence the habit of the usual sciences to avoid loops. Heinz says, ‘The hard sciences work on the soft problems and the soft sciences work on the hard problems.’ [Audience murmurs.] Quite wonderful.

HAPPY BIRTHDAY, HEINZ



Figure XXII: Heinz von Foerster in his home in Pescadero, California, in the 1990s

I’d like to move now to my two questions for Heinz’s birthday. [Figure XXII: Heinz von Foerster in his home in Pescadero, California, in the 1990s]

Question #1

When I notice what we're measuring all the time with our existing sciences—technologies of engineering, the 'hard sciences', economics, and so on—it tends to be that which is *easily* measurable and quantifiable. For example, we use Gross Domestic Product as a common measure of quality of life. I'm more interested in the Gross Complexity Product. Is our world a more complex place today than one month or one year ago, and is that a bad thing? How would we measure if it's good or bad? Is there a biological measure of **cost** that is not monetary but having to do with the time it takes from us, the stress it imposes on us, and the attention required when we try to get what we want?

This concept of '**biocost**'—stress, time, and attention—leads me to wonder, how can we take disciplines such as cybernetics—holistic and systemic disciplines, that 'put things back together', rather than 'schism' sciences—and bring them to bear on understanding and regulating our daily biocost? How can we help a world that is not in good shape?

I believe that the world is not in good shape in part because of the way we measure the state of the world. The impact of our actions is based on the metrics of the schism sciences, not holistic science. If there were a way to measure progress in terms of lowering biocost rather than in terms of monetary cost or economic cost or productivity, would that be a good thing? So my first challenge to the society of cyberneticians: What is the social design that would apply cybernetics to these hard problems, and could that bring about a different world, perhaps a more ethical world?

As Humberto Maturana says, 'We think we live as rational animals but we do not. We are emotional animals who use reason to justify our desires.' Which leads me to my first question:

What are our desires? And why would we continue to use a science that pulls things into pieces when we could use a science that puts them back together and has ethical qualities?

Question #2

My second question triggers from this figure from papers by von Foerster [von Foerster (1973). *On Constructing A Reality, op. cit.*], which has two parts. [Figure XXIII: Von Foerster's Von Foerster's representation of the nervous system as a torus]




Figure XXIII: Von Foerster's representation of the nervous system as a torus

The two parts of this figure are alternate views of the same model of the nervous system as a closed system. The left part, labeled 'a', shows the fabric of the nervous system represented as black squares and connections between. Think of the 'N' squares as neurons. The label at the top is 'syn' for 'synapses'. The left part indicates 'SS' for the sensory system impinging on the nervous system with triggers from the world. Coming out the right are the actuators of the nervous system, here labeled 'MS' for 'motor system' and response, causing actions in the environment, whose effects on the environment come back around to the sensory system to close the loop. This is one of the earliest ideas I learned from Jerry Lettvin but it took me 20 years to fully appreciate what he was talking about: the motor system completes through the environment what the sensory system initiates. Of course, it can equally well be viewed as starting from the motor system wherein the sensory system 'completes' the loop. It is simply a matter of perspective as to where you open the loop to begin the description, and so to find where it closes.

Heinz also draws a vertical loop to indicate that there are other dimensions impinging on the nervous system, the endocrine system, hormone system, immune system.

The doughnut to the right (labeled 'b') simply expresses the topology of the left in 3 dimensions, where it closes in on itself both horizontally and vertically. It's a torus.

One day I asked myself: If this is the topology of biological systems, what could it inform? I draw the following crude figure: [Figure XXIV: Reinterpreting von Foerster's torus: user and machine combined]

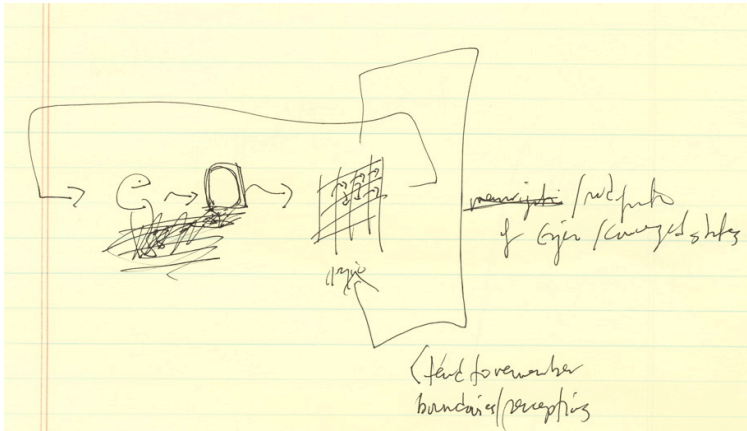


Figure XXIV: Reinterpreting von Foerster's torus: user and machine combined

The grid in the center represents the nervous system—the same representation and mechanism from the previous figure—which shows sensory input coming from its left. But here I have interposed an individual interacting with a computer screen, that in turn provides sensory input to the nervous system. Coming out of the back of the nervous system, where the motor system or effectors were in the other diagram, actions loop around and impinge on the user, this little face with an eye in it, who in turn interacts via the computer screen, completing the loop.

Also, like Heinz's diagram, there is another loop starting out of the top of the nervous system that is labeled—well, I can't read it right now, but something about 'Eigen/converged states'.

What's going on with this system? Well, I believe that this system is different than all of today's models for human-computer interaction. If we're going to take what I've been talking about—holist sciences, ethical behavior, computing a stable state, biocost—couldn't we build computer systems that had more the qualities of biological systems, and that didn't turn us into trivial machines. (Do you notice how that happens? Some new email comes in and you get the alert sounds and you immediately run to check it—you are being turned into a trivial machine.)

Cybernetics has always been interested in the creation of artifacts that are bridge the organic and the non-organic. So:

Can we make our machines part of a fabric of an extended nervous system, where its functions help us get what we want in a way that does not compromise our biology? Put another way, can we use the ideas of cybernetics to construct environments that enable us to be more who we wish to be?

That is my second question.

'GETTING WHAT WE WANT'

These two questions have a common core, bridging two domains. The first question, in the domain of the diverse cultural system, asks how we improve our well being as informed by cybernetics. The second question, in the domain of the every-day, is about how we improve minute-to-minute experience when we are using these damn stupid machines that are so far from biological, and that turn us into trivial, predictive, reactive machines.

Cybernetics, as Ashby says, is about 'getting what you want.' This is not intended in a selfish way, it is intended in the broadest possible sense. As Heinz often said, 'We see ourselves through the eyes of the other'. We see ourselves as reflected. We construct machines that reflect who we are and what we want to be. I think these machines *stink*. They tend us further in the direction of measuring progress in mechanical terms. I think the culture and the well-being of the planet are at risk. I invite your reflection based on Heinz's interests, Heinz's contribution, and the contribution of the individuals I've spotlighted today. Their work is inextricable; I don't know whether Pask said this or he said that first. *First* doesn't matter. What matters is the matrix of ideas, the matrix in Heinz's sense of the matrix that embeds, where we have the capacity to do things that we otherwise would not have. We have the capacity to create a world in which we can have wonder, in which we can be surprised with delight. And a world in which we can be the biology that we are, meshed with—not trivialized by—the complicated machines that we create. I invite these questions in honor of Heinz. Ladies and Gentleman, thank you very much for your participation with me today.

EPILOGUE

At one point during the same 1983 conference where Heinz declared himself a Paskian, a band of cyberneticians were trying to gain entrance to the conference building in the midst of a blizzard in Montreal. Each of us tried various doors. Stafford Beer, Carl Auer, Gordon Pask, Heinz, myself, and others found door after door in this block-long array of entrances to be locked. I was dispatched to find a telephone, not because I was the youngest but—at least so they led me to believe—because I was the only cybernetician in possession of both hat *and* gloves in that sub-zero weather. Stafford had his large, colorful, knit hat; Heinz had gloves; Gordon had only his cape; and so on. I trudged through the blinding snow, found a payphone, called the security office, and was told emphatically that one of those doors was certainly *not* locked; on my return and a careful, exhaustive test of all the doors, we found this to be the case. Gratefully we tumbled into the warm lobby in a pack.

With all those players present, this seemed part of my initiation into cybernetics. What mocking irony of the honest role for reductionist problem formulation in a cybernetic Weltanschauung!

My next occasion with Heinz was at the 1984 conference of the American Society for Cybernetics. There was an evening of 'oral tradition' where attendees told stories of their experiences in the cybernetic community. Heinz was a magician here, conjuring story after wonderful story of Margaret Mead, Gregory Bateson, Warren McCulloch, Norbert Wiener, and everyone he knew so well from the seminal days of our field. At one point I contributed my small tale of the blizzard and the doors in Montreal, repeating as if a mantra, 'I *am* a cybernetician'—implying how little good *that* was for getting into the building. After the event ends, here is Heinz coming toward me, grabbing my hand and shaking it vigorously. Blue eyes flashing, he looks up at me from his elfin frame and says in his persistent, energetic way, 'Paul, as a cybernetician, you stand on the Shoulders of Giants!' Bathed in the torches of his eyes, overwhelmed by his attention, how could I possibly respond? But there was no need, for he wasn't finished. Holding my gaze in his, he continues: 'Tell us....what...you...see!'

ACKNOWLEDGEMENTS

I wish to thank Karl Müller and Albert Müller and all of the many individuals behind the scenes that we don't know, for bringing about this extraordinary conference in honor of Heinz von Foerster. Especially I'd like to recognize the generosity of the family, the generosity that began with Heinz and Mai and continues through their sons to the next generations.

I'm grateful to Bernard Scott for his corrections to my live presentation concerning the timing of the development of 'conversation theory' as implied by the publication of the Pask chapter in Negroponte's *Soft Architecture Machines*. I wish also to thank Gertrud Hafner for supporting the reconstruction of my live lecture in Vienna 2003 into this form, with acknowledgement toward Heinz, who so often captured his personal magic from live lectures onto the living page. To CJ Maupin, for her unerring guidance and co-evolutionary partnership, I dedicate this paper.

SPEAKER BIOGRAPHY

Paul Pangaro studied Computer Science and Humanities at the Massachusetts Institute of Technology, US (B.S 1974), working on research contracts in the laboratories of Jerry Lettvin and Nicholas Negroponte, and he studied Cybernetics at Brunel University, UK (Ph.D. 1986), working on research contracts in the laboratory of Gordon Pask. Through the 1980s, Pangaro developed training and collaboration software applications in the company he founded, applying Pask's Conversation Theory to complex decision making applications for the US and UK governments, and the US nuclear power industry. During the early 1990s via clients such as Du Pont and Lotus Development Corporation, he built cybernetic models of information systems and human capital in the network economy. In the mid-1990s, he participated in technology startups using cybernetics as an analytical lens in the design of distributed, collaborative conversations. Pangaro was Senior Market Strategist at Sun Microsystems where he used conversational models to create strategies and relationships with software developer communities as a means of fostering co-evolution with the goal of developing new products and new markets. He is currently CTO at Snap.com. His published papers and further details on his work in user interface design, user experience philosophy, and cybernetics can be found at <http://pangaro.com>.