

Cybernetics and the Art of Living

I had intended to begin this talk, as one so often does to break the ice in plenary sessions, by being funny about some trivial matter. But Professor Trapp's sad announcement has put me in a state of shock.

Having been traveling for the last two weeks, I had not heard of Gordon Pask's death until this moment — and the loss is far too great to grasp. Gordon had generously supported me some thirty years ago when I entered academia in the United States, and there is no end to the help his ideas gave me in developing my thinking. Now I am crushed by the profound regret that a friend has gone whom I have not thanked nearly enough for the inspiration he provided.

So let me start my talk by saying I sincerely hope that a knowledgeable scholar will collect and record the history of the first ten or fifteen years of cybernetics before too many of the fascinating personal and intellectual details of its inventors are irrevocably forgotten. I, unfortunately, am anything but a knowledgeable scholar. I only had by chance some opportunities to witness the development of this revolutionary discipline as an enthusiastic outsider who was profoundly influenced by it.

The point that struck me at the outset was that the founding fathers, especially Norbert Wiener and Warren McCulloch, thought of their enterprise not merely as a technique but also as a new and powerful approach to philosophy. But the two interests quickly separated, and the spectacular mathematical and technological successes of cybernetics have until recently all but obscured the philosophical potential.

Judging by the list of symposia announced in our present program, the focus of this conference, too, is largely upon technical innovations and new applications in a variety of domains. I am therefore taking something of a risk by talking about cybernetics and the art of living. I can only hope that in the end you will forgive me.

I want to begin by recalling a statement Warren McCulloch made in a lecture at the University of Virginia in 1948. "To have proved an hypothesis false," he said, "is indeed the peak of knowledge." The "real" world does not show us when we are right, but when we are wrong. All we experience are the constraints that prevent us from acting in certain ways.

When I read McCulloch's paper in the early sixties, I had long been thoroughly dissatisfied with traditional epistemology. The statement was a revelation. A little later I came across Gregory Bateson's paper on "Cybernetic explanation" (1972), in which he explained that what makes cybernetics different from other scientific

enterprises is the fact that it operates with constraints rather than with efficient causes. He cited the theory of evolution as a prime example, because natural selection only eliminates what does *not* fit. The properties that allow an organism to survive are not created by selection but are the result of random variations.

The theory of evolution applies to species and to the heritable properties that characterize them. Species have no knowledge, they are what they are, and the organisms that compose the species either have the properties that enable them to survive, or they don't.

But there are many organisms that we call intelligent because they are able to learn from their experience. What they learn, of course, is not heritable — but it may help them to survive. They learn to avoid some constraints of the world which they experience. In other words, they learn to fit better between the obstacles their environment puts in their path.

If one takes this idea of fitting and applies to the problem of how we gain the knowledge on the basis of which we try to lead our lives, one comes to a theory of knowing that is radically different from most of the epistemologies of traditional philosophers.

Because Vienna is something of a stronghold of “Evolutionary Epistemology”, I want to stress that the model of cognition I am talking about does not fit that mold. The main reason is that knowledge in the cybernetical model is never knowledge of a real world. It is knowledge of what one can or cannot do. The obstacles that manifest themselves as constraints are merely the limits of the space that is accessible to experience. They are relative to the organism's way of experiencing, not representations of an independent reality. Knowledge, in this theory, is therefore not a picture of reality, but a repertoire of actions and thoughts which in past experience have turned out to be successful.

In this sense, this theory of knowing replaces the notion of *true* representation with the notion of *viability*. — Rather than go into the details of that theory which is laid out elsewhere (cf. Glasersfeld, 1995), I shall give you tangible examples of how I see it.

I spent the last ten days in Chamonix and the mountains around Montblanc. It was a nostalgic experience, because until forty years ago I spent many a spring skiing on the glaciers of the Alps. In those days there were no cable cars and other mechanical devices to bring thousands of skiers to the tops of mountains. You were alone there, and if you wanted to ski down a mountain, you first had to climb it.

In retrospect, it struck me as a good example of dealing with constraints. If you wanted to go up or down a mountain, you had to look at it rather carefully. You wanted to reach the summit — but it would have been a mistake simply to look for an easy way up. As an experienced mountaineer, you first of all figure out where you must *not* go. You try to see possible avalanches, ice breaks, crevasses, and other fatal constraints. Only when you have, so to speak, blocked out the treacherous parts of the mountain, would you begin to plan your way up. At this point, you do make choices, but you make them *within* the space left between the mountains constraints. To “know” a mountain means to know where, on its slopes, you are relatively safe; it means to have learned the viable paths.

One can come to the notion of viability in many ways. One of them is the principle that Leibniz and Maupertuis formulated a long time ago: The principle of least action or, respectively, of least resistance.

Water will follow the pull of gravity as far as it can. When it rains on a hill, the rain water runs down wherever it finds a way. If it is stopped, it will collect and eventually flow over or around the obstacle. This, in turn, changes the shape of the hill, makes new paths viable, and encounters new constraints.

But let me return to the origins of cybernetics. Besides shifting the focus of attention from causes to constraints, it brought about another fundamental change by launching the theory of communication.

When Claude Shannon formulated the theory mathematically, he was careful to state in the first two pages of his seminal paper that what he called “information” had nothing to do with semantics. The impulses that travel in a communication channel from a source to a sink are changes of some form of energy. They are “signals” only to those who are in possession of the relevant code. The code itself is not part of the transmission. The signals are instructions to select specific parts of the code. And “information”, in the theory of communication, does not refer to the meaning of the coded elements, but is simply the measure of how many or how few of the pre-established elements the signals select.

You probably know all this — but it is good to remember it when someone is speaking.

Norbert Wiener provided the marvelous example of the flower shops that use their own economical system of communication. If a young man in Vienna spent a few happy days with an American tourist and, now that she has left, he wants to deepen the impression he made on her, he might go to a flower shop and arrange for a dozen red roses to be sent to her in Los Angeles on her forthcoming birthday. The flower shop then cables the address, the sender’s name, and a specific number, say 54. By means of two simple digits the number instructs the receiving shop to select 12 red roses and the message “Happy Birthday”.

Don’t think that I have done this so often that I know the florists’ code by heart. I just invented the number 54. But florists do use such numbers for the selection of specific flowers and all sorts of good wishes and condolences. My point in mentioning it here is simply that the number is meaningless *unless* it is interpreted or decoded by a person who knows the code. The “information” the number carries in this context is no more than that it indicates and thus selects a particular item out of all the messages the florists’ code contains.

This fundamental condition was unfortunately disregarded by nearly all the linguists who went into a frenzy of excitement when they heard of Shannon’s theory, and they promptly compounded their confusion by speaking of “Information Theory”, while Shannon had deliberately called it “Theory of Communication”.

Human language is, of course, a communication system and it is therefore quire enlightening to apply Shannon’s theory to it. But language is also different from all artificial or technical communication systems. The crucial difference is that in language we do not start out with a pre-established code, but each of us to learn it by using it.

You may say this is nonsense, because we have dictionaries that tell us the meanings of words. True, we have dictionaries — but how do they tell us the meanings of the words we look up? They use other words. Just think for a moment how far you would get if you had to learn Morse code by trying to use it.

A one-year old child, in the process of acquiring language is almost in as difficult a position. I say “almost”, because there is an important difference: the language the child has to learn is constantly being used by the speakers within the child’s field of immediate experience. An example may help to show what I mean. Let us say a mother says to her child: “It’s time for your bottle.” She goes to fetch the bottle, puts the nipple to the child’s mouth, and says: “Drink your milk!”

The child begins to suck (because it sucks everything that touches its lips) and it feels the liquid in its mouth. No doubt it will form some association between the sound of the words, the touch of the nipple on its lips, and the feeling in its mouth. But it will take a great many other experiences with drinking water or orange juice, with cups and glasses, and with many other word-sounds, before the child has sorted out approximate meanings for “bottle”, “drinking”, and “milk”. And the most important aspect of this learning situation is this: the experiences of bottle, drinking, and milk, with which the child associates the sounds of these words, are the child’s subjective impressions. They are neither the mother’s nor anyone else’s. Nor are they “things in themselves” or instances of independent objects in a real world. They cannot be anything but the impressions this child-subject happens to experience.

Clearly, in the case of words that are frequently used in everyday language, these subjective impressions become more or less intersubjective in the course of linguistic interactions with other speakers. But one can show that even the commonest words retain a margin of subjective meaning for each individual speaker. Linguists and philosophers of language usually subsume this margin under the term “connotation”, and they claim to be able to separate the subjective component neatly from “objective denotation”. From our point of view, this claim rests on the illusion that words *refer* to things in a real world. In our theory — which we of course consider to be more adequate — words, as I hope to have shown, refer to subjective experiences of the individual language user. The separation between denotation and connotation thus no longer involves objectivity, but becomes a question of greater or lesser *fit* with the usage of other speakers.

The result of our investigations in this area is that the meanings of words and longer segments of language are never “shared” with others in the sense that they could be considered the same for all members of a language community. All one can say is that among proficient speakers of a language, meanings are at best *compatible*, i.e., they function similarly in most situations.

This is an important difference from the artificial, technical communication systems. There, the code that bestows significance to the signals is established and distributed to the communicators *before* any communication takes place. In contrast, a child acquiring its human language enters into a system that is already in action and as newcomer it has to establish a code for him- or herself. This is a laborious process that involves countless trials and errors and leads at best to viable approximations. In fact, it is an endless process. No matter how old you are and how long you have been speaking your language, every now and then you discover that you have been using a

particular word in a way that is not compatible with the accepted usage of your community. You were unaware of the idiosyncrasy, simply because the situation in which the discrepancy becomes relevant had not occurred in your past experience.

On the surface, this may seem to have little to do with the art of living, but it is obviously a factor in the art of reading. How often do you open a book on an intellectual subject, and in the first few pages you come across a statement that seems quite nonsensical. If you are an impatient reader, you may say to yourself, this author is a fool — and you put the book away. But if, instead, you keep in mind that the meanings of words are essentially subjective constructs of individual language users, you will tend to withhold judgement. The author, you will say to yourself, is supposed to be intelligent and therefore it is likely that what he or she has written makes sense to the writer. In this case you will make an effort to find out what this sense could be. Quite often such an effort is worth making, because it may lead to the realization that the text not only uses some terms in an unfamiliar way, but also that this new way makes good sense. Whenever this is what you manage to find, you have learned something new — and that, after all, is the deeper purpose of reading.

How many philosophical debates could be turned into productive discussions, if only the participants were not quite so convinced that the meanings *they* have associated with words are the only legitimate ones.

How many quarrels between lovers could be avoided, if one of the two considered that what the other says may not mean what it appears to mean.

All this, of course, raises the question of what *we* mean when we say we have understood a piece of language. There is still the wide-spread view that words *contain* their meaning the way a book contains pages. If, however, words are printed or spoken signals in the linguistic communication system, they cannot convey a fixed meaning. They can only point to and select whatever the reader or listener has associated with them. And the conceptual structures that this person has associated with the given words are abstractions from that individual person's experience, not from the experience of the writer or speaker. No doubt each language user's associations have been adapted and honed by years of linguistic interactions with others, but the material of which they consist is under all circumstances subjective experience.

What, then, is understanding?

I want to suggest that understanding depends on the sense you make of what is said or written. If the concepts the words have called up in you, and the way the sentences have prompted you to relate them, yield a conceptual network that fits the context created by what came before and is not countermanded by anything the speaker says or does now, then you assume that you have understood what he or she intended.

This, of course, is a simplification. What I called context is usually a hierarchy of different contextual levels, such as your past experiences with the speaker or author, your own construction of the experiential world in general, certain expectations you have formed, and other things as well. However, the point I want to emphasize is that on all levels it is a question of *fit* — not a question of receiving or reproducing conceptual structures that originated in the speaker's head. What a

speaker or author wants to say is forever inaccessible — you can only interpret what he or she actually said.

You may wonder why I spent so much time talking about language. I had two reasons. The first is that much of our living is done in conjunction with others, and language is inseparable from the social context. I therefore feel that a coherent model of how linguistic communication works is a great help in managing our social interactions and thus our life.

It would make our lives more pleasant, for instance, if every time one is about to shout at someone: “But I told you so!” one remembered that telling does not guarantee being understood.

The second reason is that the principle of interpretation and viable fit that I have outlined with regard to understanding language is equally applicable to understanding the world in which we find ourselves living. We have no more access to an ontological reality than to the thoughts of another person. All we have to go on is our experience. In both cases we interpret what we see, hear, and feel, and we construct models that should enable us to make predictions.

At an earlier edition of this conference, some years ago, I suggested that if the model we have constructed of the person we live with has served us well for some time, we tend to believe that it has captured how that person really is. But sooner or later our companion does something that we did not expect. This may irritate us, and we reproachfully say: “You have changed!” — Often this is not at all the case. The other has merely shown an aspect we had not incorporated in our model, because no prior situation has brought it to the fore. Our surprise and our irritation would be greatly mitigated if we kept in mind that the other we know is not the other as he or she is, but a model we have constructed on the basis of nothing but our own experience.

The very same happens to the scientist who constructs a model of, say, the planetary system or the universe. If that model works well and provides useful answers to the questions that are asked, it comes to be regarded as a true description of reality. But sooner or later something incompatible is observed, a precession of Mercury or a beam of light that does not follow a straight line. Such observations constitute constraints that demolish the viability of the accepted model. It no longer fits the scientists’ expanded experiential world. At first there usually is some incredulity and a great deal of reluctance to accept such failure. Eventually, however, a new model is constructed with the help of new concepts that make the shocking observations seem normal and expected.

All this is in harmony with the fundamental principles of our discipline, for cybernetics is the art of creating equilibrium in a world of possibilities and constraints. — And I would suggest that this is also a viable definition of the art of living.

References

- Bateson, G. (1972) Cybernetic explanation. In Bateson, *Steps to an ecology of mind*, (pp.399-410). New York: Ballantine.
- Glaserfeld, E.von (1995) *Radical constructivism: A way of knowing and learning*. London: Falmer Press.

McCulloch, W.S. (1948) Through the den of the metaphysician. Printed in McCulloch, Embodiments of mind, (pp.142-156) Cambridge (Massachusetts), M.I.T. Press, 1965.

Shannon, C.E. (1948) The mathematical theory of communication. Bell Systems Technical Journal, 27, 379-423 & 623-656.

This paper was downloaded from the Ernst von Glasersfeld Homepage, maintained by Alexander Riegler.



It is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/2.0/> or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, CA 94305, USA.

Preprint version of 16 June 2012