

H. VON FOERSTER

**BIOLOGICAL PRINCIPLES OF INFORMATION
STORAGE AND RETRIEVAL**

BCL

P U B L I C A T I O N

150

REPRINTED FROM

ELECTRONIC HANDLING OF INFORMATION:
TESTING AND EVALUATION, Kent, Taubee,
Beltzer, and Goldstein (ed.), Academic
Press, London (1967), pp. 123-147

BIOLOGICAL COMPUTER LABORATORY
Department of Electrical Engineering, University of Illinois, Urbana, Illinois

*BIOLOGICAL PRINCIPLES OF
INFORMATION STORAGE AND
RETRIEVAL**

by
Heinz Von Foerster
University of Illinois

Dr. Ruth Davis, in her keynote address (Chapter I), tossed out to all of us a very important challenge. If I translate this challenge into my own language, then essentially she said that all the large information handling organizations are, de facto, functional isomorphs of living organisms. And since this is so why don't we, in our analysis and synthesis of these large organizational systems, look first at the structure and the functions of these living organisms.

Of course she didn't use exactly the same words as I did right now. For instance, she didn't say "living organisms" she just simply said "people." Or, for instance she didn't say "functional isomorphs," she just said

*Some of the ideas presented here grew out of work currently sponsored by AF-OSR Grant 7-67, AF Contract 33(615)-3890 and AF Contract 49(638)-1680.

"duplicative mimicking." But essentially I think she meant exactly that because she illustrated with her series of examples very beautifully precisely that point.

If I remember correctly, as one example she cited the large military organization, another one was a business concern and the third one was, I think, a large data retrieval system.

However, she showed that in these organizations there still remain a lot of tenacious bottlenecks which apparently are always present. They don't disappear. They're just shifted around, pushed from one side to the other side, and the load of the responsibilities are just transferred from one shoulder to another, but never resolved. On the other hand, if one looks at biological systems, how they handle information, one may see that some of these bottlenecks are not at all present. All responsibilities are beautifully distributed and the system functions very smoothly. With her challenge to consider structure and function of biological systems and the way in which they handle information, I think she put her finger on an extremely important point in the whole field of information storage and retrieval. Of course at the same time she raised an extraordinarily difficult and very complex issue.

Nevertheless, it is precisely to this issue that I would like to address myself this morning. Unfortunately, I can do this only in a very superficial fashion, hoping that at least you get a feeling of what we believe is going on in rather complex living organisms and of the way in which they handle information, for these systems handle information in an entirely different way than we do in our information storage and retrieval systems. Moreover, they process the information fast, effectively and reliably.

Perhaps you may get already a hunch of the fundamental difference between the way in which living organisms handle information and the ways in which we do this now by contemplating for a moment the fact that living organisms perform better and more efficiently with increasing complexity of their tasks, while of course, our information systems at the moment become increasingly slow and hopelessly deficient if they are confronted with tasks of even moderate complexity.

For example, we all know that a three-year-old child easily speaks a natural language; a three-year-old child recognizes its makers in about one hundred milliseconds; however, a three-year-old child has great difficulty in multiplying 269 by 3,625. However, a computer can do this very fast, as you know. This is a very trivial task and hence it does it in a jiffy. But of course, he has tremendous difficulties and he's very clumsy

indeed if he has to communicate with us in natural language. In addition, I have never seen a computer recognizing its maker. Have you seen a computer wiggling its tail when the head of the computer department enters the lab? Clearly, there seem to be some fundamental differences.

The problem to which I would like to address myself now is the basic issue in all the information processing of living organisms and I shall attempt to develop the various features of information processing in living organisms from a single first principle. This first principle is essentially the principle of inductive inference, as you will see in a moment. But let me build up to that point. I just mentioned it so that you know what my intentions are.

You know that the essential problem of a living organism is not so much to live. It is essentially to *sustain* life, it is the problem of survival. That is to live for the next instant also, not to live only for the present moment and then to perish at the next one. In other words, what this organism has to do is to see that it never gets itself into a situation so that it cannot get out of it again. It has, so to say, to anticipate future events in order to avoid the possibility, for instance, of being crushed by some unfortunate circumstances. It can do this only, of course, by figuring out from past sequences of events future sequences of events. But precisely this computation of future sequences of events based on the experience of past sequences of events is, of course, the principle of inductive inference.

This principle of inductive inference sounds in this formulation like sheer magic. That means, I seem to propose that every living organism is, so to say, a little vest-pocket soothsayer who predicts what will go on in the future. Of course, the principle of inductive inference does not claim to be such a thing at all. In fact, the temporal aspects of this principle are not at all relevant: it is essentially a principle of generalization, as you probably remember. It simply says that of all things, x , which have the property P_1 and have the property P_2 , that all other things that are not x and have the property P_1 will also have the property P_2 . Essentially what one does here is to establish concomitances of properties and make this concomitance the general property. Clearly, this is the inverse operation of deduction where one derives from a set of axioms all the conclusions. In induction one starts with the specifics and projects into the generalities.

For those of you who may appreciate more rigor in discussing induction, let me present in passing this principle in Russell-Carnap notation. They use the following symbolism:

$(x) \dots$ "for all x "

$(Ex) \dots$ "there exists at least one x "

$P(x) \dots$ " x has the property P "

$\& \dots$ "and"

$\rightarrow \dots$ "implies"

$\neq (x,y) \dots$ " x is different from y ".

With these, induction is formulated as follows:

$(x) (Ey) \{ P_1(x) \& P_2(x) \rightarrow [\neq (x,y) \rightarrow P_1(y)] \} \rightarrow (y) [P_1(y) \& P_2(y)]$

Or in other words: "if for all x for which P_1 and P_2 are concomitant there exists a y , different from x , with property P_1 , then—and now comes the big jump to generality— P_1 and P_2 are also concomitant for all y 's."

Clearly, this statement is not always true, irrespective of its propositions being true or else false; that is, this statement is not a tautology, hence it must say something about the world and thus can be found true or else false. But this is exactly the fate of all theory: they may last for a while until some observation is made which contradicts the hypothesis. "Back to the drawing board!" is the accepted response among scientists whose essential job is to present their hypotheses so that, in principle, they can be shown wrong. A hypothesis that cannot be shown wrong says absolutely nothing as, for instance, the hypothesis: "Tonight I shall, or I shall not, go to the movies."

Since the principle of induction is not always true one is tempted to justify this universally applied principle. However, the first to discover that we would run into peculiar difficulties if we'd like to justify this principle was, I believe, Hume. He pointed out that if we try to justify the principle of inductive inference by saying, "Look, people, it always worked in the past, and therefore, why shouldn't it work in the future?" we use the principle of induction to prove the principle of induction. And that is clearly not legal, to say it politely.

A most delightful and penetrating discussion of this problem has been given by Jerold Katz (1). I can only recommend this treatise to everyone who enjoys a brilliant argument. However, one may assume a much tougher position regarding this principle which is the basis for our concept of causality. Ludwig Wittgenstein (2) submitted this concept to a severe criticism. In his opinion "the belief in causality is superstition," a statement that comes out in its original German much more beautiful: "Der Glaube an die Kausalität ist der Aberglaube." He insisted "That the sun rises tomorrow is a hypothesis. In fact, it may not rise tomorrow." And he is so right. Indeed it may not rise tomorrow.

With this extravaganza I wanted just to illustrate that we are in a somewhat peculiar situation when we rely on this principle. And we rely on it, otherwise we would not be alive. This fact suggests to Wittgenstein that this principle implicitly expresses something about the world, namely that concepts like "causality" are formulable. If I may translate that for my own purpose, I would say that it shows that there is an intrinsic structure in the world, because if it were not structured, no descriptions of the universe could be made.

At this point you may now get a little nervous. Here I come and tell you about the principle of inductive inference while this is a conference of information storage and retrieval. What you apparently came here for is to learn about theories, methods and systems which permit you to retrieve from a large file of documents some pertinent data in a hurry. What has this to do with inductive inference?

To see that these two points are closely related let me first make clear that in almost all cases of inquiry one is interested in the *facts* and not so much in the *representation* of facts. Say, you wish to know which year was Napoleon's birth date. I believe that you could not care less whether the system answers with "1769," with "seventeensixty-nine" or, perhaps, even with "MDCCLXIX," as long as you know it was seven years before the Declaration of Independence. (Note the various concomitances that can be established.) However, you may argue that this is a trivial example, although I do not believe it so trivial at all. Nevertheless, let's pose a more complex problem to our system. We wish to know, say, how the famous Sea Battle at Midway in World War II evolved, how the ships moved, how they retreated and advanced, and how many ships were incapacitated or sunk. This would be one way of posing the question. Another way would be to ask how Admiral Nimitz reported the development of this battle.

From this it appears that in the former case one is after the facts, while in the latter case one is after a representation of facts. Not so, I venture to say. For in the latter case we will learn from the way how Admiral Nimitz describes this battle more about Nimitz than about the battle. To know something about Nimitz is again knowing facts. One may learn more about the battle from an account of a Third Mate on a little gun boat, an account that may be quite different from that of the Admiral, and that for quite obvious reasons. As you may remember, the flagship was the first that was knocked out and Nimitz did not know at all what was going on during the battle. But to repeat: In the way in which he reconstructed the battle, you will learn something about him. That is,

Now, if we are really after the facts, and if we are only secondarily interested in the description of facts, then it is not at all necessary that we have to retrieve word by word, comma by comma, period by period of an arbitrary document that gives one of a million different ways to account for the fact. Hence, there is no need to store all that junk. It is completely sufficient to store some abstract structural—not necessarily linguistic—representation of facts by computing the general principles of what was going on and, upon inquiry, to regenerate from these, by using linguistic constraints, all that can in fact be known about the facts we wish to know.

With paraphrasing the information storage and retrieval problem in this way, I hope that I have shown its close relation to the problem of inductive and deductive inference. To store is to generalize by induction, and to retrieve is to specify by deduction.

At first glance this may appear unnecessarily complicated and one may argue for the simple, straightforward method to store all junk and to leave it up to the retriever to select what appears pertinent to him at that moment. I hope I will be able to show in a moment that this attitude is wishful thinking if we consider files that are not kept small or have an exceedingly simple internal organization. In fact, living organisms could not survive, if they would incorporate all sensory data into their "storage and retrieval system." First of all, there would be not enough storage space available; second, this hodgepodge of data would not lead to any insight into what the world is all about. But I shall return to this point later.

At this moment, however, I shall try to clue you in on the "what" and the "how" of information storage and retrieval in an inductive and deductive inference computer by resorting to a metaphor.

Assume that I am confronted with a task that requires a large number of multiplications that have to be done with great reliability. Since I am—I am sorry to say—very weak in multiplication (rarely get the right result) I may solve this problem for myself once and for all by preparing one of the standard multiplication tables with two entries X as rows, Y as columns, and the product $X \cdot Y$ in the body of Table I. In preparing this table I wanted to know how much paper I need to accommodate factors X , Y up to a magnitude of, say, n decimal digits. Using regular-size type for the numbers, on double-bond sheets of $8\frac{1}{2} \times 11$ inches, the thickness L of the book containing my multiplication table for numbers up to n decimal digits turns out to be approximately

$$L = n \cdot 10^{2n-6} \text{ cm.}$$

TABLE I

X·Y	Y								
	0	1	2	3	4	5	6	7	. . .
0	0	0	0	0	0	0	0	0	. . .
1	0	1	2	3	4	5	6	7	. . .
2	0	2	4	6	8	10	12	14	. . .
3	0	3	6	9	12	15	18	21	. . .
4	0	4	8	12	16	20	24	28	. . .
5	0	5	10	15	20	25	30	35	. . .
6	0	6	12	18	24	30	36	42	. . .
7	0	7	14	21	28	35	42	49	. . .
.	.	.	.	.	.	.	.	.	. . .
.	.	.	.	.	.	.	.	.	. . .

For example, a 100×100 multiplication table ($100 = 10^2$; $n = 2$) fills a "book" with thickness

$$L = 2 \cdot 10^{4.6} = 2 \cdot 10^2 = 0.02 \text{ cm} = 0.2 \text{ mm.}$$

In other words, this table can be printed on a single sheet of paper.

Now I propose to extend this table to 4,6,8 and 10 digits of the factors X and Y. With my formula for L of the thickness of the book, or the length of the shelves that store these books, I calculated Table II.

TABLE II

n	L
2	$2 \cdot 10^{-2} \text{ cm}$
4	$4 \cdot 10^2 \text{ cm}$
6	$6 \cdot 10^6 \text{ cm}$
8	$8 \cdot 10^{10} \text{ cm}$
10	$1 \cdot 10^{15} \text{ cm}$

In other words, if I wish to multiply eight digit numbers, my multiplication table must be accommodated on a bookshelf twice the distance from the moon to the earth, and if I wish to multiply ten digit numbers it must be about 100 times the distance from the sun to the earth, or about one light day long. A librarian, moving with the velocity of light, will, on the average, require one half day to look up a single entry in the body of this table.

With this example I persuaded, at least myself, that the simple, straightforward method of storage and retrieval by fabricating a record of all data to be potentially retrieved is not a very practical solution, and I had to look around to find a better way to do this. Since my requests were very moderate, I restricted myself to eight digits, I hit upon a little gadget in the form of a small cylinder $2\frac{1}{2}$ inches in diameter and less than 4 inches high which contains 16 little wheels, each with numbers from zero to nine printed on them (See Figure 1). These wheels are sitting on an axle and are coupled to each other by teeth and pegs in an ingenious way so that, when a crank is turned an appropriate number of times, the desired result of a multiplication can be read off the wheels through a window. The whole gadget is very cheap indeed and, on the average, it will require only 30 turns of the crank to reach all desired results of a multiplication involving two eight-digit numbers.

The answer to the question of whether I should "store" the information of a $10^8 \times 10^8$ multiplication table in the form of a $8\frac{1}{2} \times 11$ inch book one half million miles thick, or in the form of a small manual computer, is quite obvious, I think. However, it may be argued that the computer does not "store" this information but calculates each problem in a separate set of operations. I am sorry to disappoint those who believe in the magic powers of these gadgets. In fact they do nothing of this sort. My turning of the crank does nothing but give the computer the "address" of the result, which I retrieve at once—without the "computer" doing anything—by reading off the final position of the wheels. However, if I can retrieve this information, it must have been put into the system before. But how? Quite obviously, the information is stored in the computer in a structural fashion. In the way in which the wheels interact, in cutting notches and attaching pegs, all the information for reaching the right number has been laid down in its construction code, or, to put it biologically, in its genetic code.

If I am now asked to construct an information retrieval system or, if you wish, a "brain" capable of similar, or even more complicated stunts, I would rather think in terms of a small and compact computing device instead of considering tabulation methods which tend to get out of hand quickly.

It has lately become increasingly clear that evolution, in which an ever so slight advantage is immediately associated with a considerable pay-off in survival value, has gone exactly the way I just suggested, trading cumbersome accumulation of isolated operational entities for structural sophistication. Homeostasis is a typical case in question. However,




Figure 1. A MANUAL DIGITAL COMPUTER ACCOMMODATING PRODUCTS UP TO $10^{16}-1$.

since our topic in this conference is information storage and retrieval, let me give you a very brief account on some of the basic principles that, we believe, govern all processing of information in biological systems.

The usual difficulty encountered, not only in describing, but also in comprehending these principles is that it involves the appreciation of

two kinds of mapping. One that allows for an internal representation of environmental features. These are represented in the structure and function of the information processing apparatus, in principle the central nervous system. The other kind of mapping concerns a representation of these internally mapped features in the form of messages composed of symbols that are remappable into the internal representations of a partner who participates in a symbolic discourse. Let me illuminate this situation with two illustrations.

Figure 2 is a schematic of the canon of perception and cognition. Since our universe is not Chaos, that is that anything can happen that could logically happen, but Cosmos where only certain things can happen according to the "Laws of Nature" that are so beautifully, but naively, presented in textbooks of physics, chemistry, etc., "cosmic constraints" must prohibit these pages you hold in your hands to turn suddenly into big, pink elephants that fly out through your window. These constraints manifest themselves in spatio-temporal structures, the "features" of our environment (upper box in Figure 2). However, an organism looking at this environment cannot see the constraints. He only gets information about the environmental structures via his sensory apparatus. But since

ENVIRONMENT




Figure 2. CANON OF PERCEPTION AND COGNITION

he has to make inductive inferences of what his environment will be later on, or beyond what he immediately perceives, this information is completely worthless unless—and here is the important point—he can compute the constraints that produced these structures. Hence, he shoves this information to his brain and lets it figure out the constraints. These, in turn are tested in the environment and new structural information is perceived. The terminology in this closed information loop is “perception” for reception of structural information, and “cognition” for this computation of constraints. I shall in a moment discuss a simple example of structure and function of a constraint computer that maps the environmental feature “containing discriminable things” into the corresponding internal representation. First, however, let me briefly discuss the other kind of mapping.

Figure 3 is a schematic of the canon of symbolic discourse. The two partners in this discourse are the two squares (sorry, no pun) who are both exposed to the same environment consisting of a set of objects Ω , one in the process of perceiving (mapping of the previous kind, denoted by ϕ), the other in the process of testing (inverse of ϕ , denoted by




Figure 3 CANON OF SYMBOLIC DISCOURSE

ϕ^{-1}), and who communicate with each other via Σ_M , a structured set of symbols called a message. "Comprehension" of a message is a successful mapping, denoted by ψ , of this structured set Σ_M by the recipient into another structured set, a set of relations Σ_R , which is, of course, nothing else but the internal representation of environmental constraints. The inverse of this mapping denoted by ψ^{-1} is the attempt by one partner to externalize Σ_R in the form of Σ_M so that $\psi^{-1}\psi$ will indeed regenerate the appropriate Σ_R in his partner.

With due apologies for this lengthy discussion of a highly abstract and apparently unduly complex skeleton, I hasten to surround this skeleton now with concrete flesh. First, I shall briefly discuss the physiological apparatus that realizes the mapping ϕ and the storage of Σ_R , second I shall touch on some of the peculiar features of the linguistic constraints that manifest themselves in Σ_M . We will see later that the code in Σ_M cannot be cracked unless there exists the mapping ψ which permits a unique interpretation of messages in Σ_M . In other words, linguistics alone does not give us a clue about the function of language, unless it takes cognitive processes into consideration.

Cognition and perception is preceded by sensation. Sensation is mediated through elementary sensory receptors, the cones and rods in the retina, the hair cells on the basilar membrane in the ear, the receptors of touch, etc. These sensory neurons translate highly specified physical modalities as light, motion, pressure, etc., into one universal language: the transmission of sequences of electric pulses. Figure 4 shows such pulse sequences as they were measured by placing a tiny electrode in the vicinity of the fiber that emerges from a cutaneous touch receptor, and stimulating this receptor with increasing levels of pressure. The receptor responds to these increasing stimulus intensities with the transmission of increasing frequencies of electric pulses. Radio engineers call this Frequency Modulation, or FM for short, and are very proud of having discovered this means of almost noiseless signal transmission a few decades ago. Nature did it a few hundred million years ago. Moreover, physical intensities are mapped into frequencies that are proportional to the logarithm of the intensity, translating external multiplications into internal addition, for — as you may remember —

$$\log (A \cdot B) = \log (A) + \log (B).$$

Figure 5 shows an isolated neuron from the cortex of a cat. The actual size of this element in its entire extension is not larger than about




Figure 4. ELECTRICAL PULSE ACTIVITY MEASURED WITH MICRO PROBE
IN THE AXON OF A SENSORY RECEPTOR.

10 thousandths of an inch, but it represents in its function a computer of considerable complexity. Its inputs are hundreds, sometimes thousands of "axons," i.e., the fibers emerging from other neurons. These axons




Figure 5. AN ISOLATED CORTICAL NEURON.

terminate on the "dendrites," the branchlike ramifications stretching upwards and sideways. Closer inspection of this figure shows the many points of termination as tiny protuberances along these branches. The effect of pulses which travel down the axons from other neurons and reach this neuron are twofold. They may either stimulate this neuron to generate a pulse that will, in turn, be sent forth over its axon to other neurons, or they may inhibit such action which would have taken place in absence of this inhibition. Although the mechanisms of excitation and inhibition are not completely understood, in their most simple form they may be interpreted as affirmation and negation in a logical operation, the neuron representing the computer for logical functions. Figure 6 is a symbolic representation of this situation, introduced by McCulloch(3), the knob attached to the cell body signifying facilitation, the loop around the spike on the top inhibition.

This symbolism permits one to gain a quick insight into the computational property of neural networks of moderate complexity. These are usually found following immediately the layer of sensory receptors as

INHIBITION

EXCITATION




Figure 6. SYMBOLIC REPRESENTATION OF THE FUNCTIONS OF A NEURON.

may be seen in the semi-schematic sketch of the post-retinal network of a mammalian eye (Figure 7). At right a single network of three rods playing on one bipolar that acts on one ganglion cell is isolated for clarification. What are its functions?

One important function may be discussed with the aid of a simplified one-dimensional model (Figure 8). Here the sensors are located in




Figure 7. SEMI-SCHEMATIC CROSS-SECTION THROUGH THE POSTERIOR RETINAL NETWORKS OF A MAMMALIAN EYE. LAYER (1) CONES AND RODS; (2) CELL BODIES OF CONES AND RODS; (3) CONE AND ROD AXONS TERMINATING ON BI-POLAR DENDRITES; (4) BI-POLAR CELL BODIES; (5) BI-POLAR AXONS TERMINATING ON GANGLION DENDRITES; (6) GANGLION CELLS; (7) OPTIC NERVE.




Figure 8. (a) MODEL OF THE STRUCTURE AND (b) FUNCTION OF A SPECIAL NETWORK IN A ONE DIMENSIONAL "RETINA."

the top row, sending four fibers to cells into a second layer where they make contact with the "computer neurons." Consider that the two fibers from each sensor that contact the two lateral neurons inhibit, while the two fibers that lead to neurons directly under each sensor facilitate.

Under these conditions a uniform illumination (Figure 8a) will elicit no response of the whole network, independent of the strength of illumination, or any temporal variation. This is due to the double inhibition acting on each computer neuron and cancelling the double facilitation coming from straight ahead.

If, however, an obstruction is placed into the light path (Figure 4b), the edge of this obstruction will be detected at once, because the only neurons which will now respond are the ones on the edges of the obstruction, receiving insufficient inhibition from only one of its neighbor cells in the light, while the other one is in the shade and silent.

It may be worthwhile to note that this network can detect a property which cannot be detected by the human nervous system. Consider the simple topological fact that any finite, one-dimensional obstruction must have two edges (Figure 8b). If n objects obstruct the light path to this "edge detector," $2n$ neurons will be active, and their total output divided by two gives exactly the number of objects in the visual field of this one-dimensional "retina."

I went into these details of computation in a particular model of a neural net not so much to show you a nice "edge-detector" or an amusingly simple counter. I wished to demonstrate a basic principle of neural computation, namely, the computation of abstracts. By this I mean that if, say, three objects are in the visual field of this "retina" the network will report the presence of three objects independent of their location, their movements, their size, and independent of the intensity of illumination beyond a certain threshold. In other words, structure and function of this network "abstracts" from all these accidentals a particular property of its environment, namely the presence of three distinct objects. It represents, if you wish, the Platonic Idea of "distinct, separable entities."

Networks having other structural and functional properties compute other abstracts, and today we have a fair knowledge of the relationships between these two (4). Permitting variations in these structural and functional properties, these networks may learn to compute a large variety of abstracts which, in their almost infinite variations and combinations, can "hold" most anything that is desired (5). Clearly, these networks cannot keep a record of the events that were presented to the sensors. However, they can hold the relations of abstracts and thus have all the necessary and sufficient prerequisites for allowing the system to draw inductive inferences (6)

I hope that despite this admittedly superficial treatment of some of the esoteric points in neurophysiology I have given you sufficient clues

as to the mechanisms that are involved in what I earlier called the "mapping ϕ ". I shall now turn to the other kind, the "mapping ψ ", which will prove to be the only way in which the meaning of an otherwise ambiguous, multivalued "message" will be uniquely determined.

In order to see to which problem I am addressing myself now, let me briefly recapitulate some of the peculiar properties of language *per se*. The first of these properties is that language is a closed system. By this I mean that although language refers almost always to extra-linguistic entities—"this table," "that person," etc.—when it is robbed of extra-linguistic behavior as, e.g., denotative gestures, ostensive definitions, etc., it is impossible to learn a language through language, for any word in a language is again defined by words in the same language. If all explaining, and to-be-explained, words are unknown, they must remain unknown. A dictionary of a single language is a typical example of circular definitions in language.

I give an instance which, with amusement, may be extended to any other case. Suppose I do not know the meaning of "after." I consult a dictionary which explains "after" with the aid of "later," "subsequent," "succeeding" and "consecutively." Suppose I do not know the meaning of these. Consequently I look them up in my dictionary and again I do not know the terms given now, and so forth. The result of this search is given in matrix form in Table III which lists for the American Collegiate Dictionary on the left hand side the source words, on the top the target words, and indicates with an "X" at intersections of rows and columns the presence of a reference.

Clearly, if I do not know any of the six terms, I shall never know what "after" means. How to escape this dilemma? It has been suggested that denoting by pointing is a meta-language, that breaks through the closure of language. Alas, this meta-language may also be closed, as told in a charming story by Margaret Mead, the anthropologist. She wanted to learn quickly a basic vocabulary of a people whose language was unknown, and she pointed at various things with questioning sounds and gestures. However hard she tried, she always got the same utterance as answers. Later she learned that these people were helpful, concise, and understanding: the utterance was their word for "pointing a finger."

A second property of language I wish to mention briefly is that, at least for Indo-European language, a major subsystem, the nouns, are ordered hierarchically. By this I mean that a noun is usually defined by another noun of large generality to which specifiers are attached. For example, looking up "bee" in a dictionary one finds

TABLE III

Illustration to the assertion that
language is a closed system.

	AFTER	LATER	SUBSEQUENT	SUCCEEDING	CONSECUTIVELY	FOLLOWING
AFTER		X	X	X	X	
LATER	X				X	X
SUBSEQUENT						X
SUCCEEDING	X		X	X		X
CONSECUTIVELY						X
FOLLOWING	X					

"any of a number of related four winged, hairy *insects* which feed on the nectar of flowers."

To play the game of before, one may look up "insect" and will find it to be an *animal* of certain specification, etc. Ultimately, however, one will arrive at a single, most general noun which, depending on the initial entry, may be "thing," "state," "action," etc., which is not further explained, or if "explained," then in the sense of circular definitions as in the previous case. We consider such terminal nouns as the "roots" of a tree from which the various nouns emerge as its branches. A short segment of such a tree is given in Figure 9 which is obtained by entering a dictionary at random, picking a noun and constructing from this initial entry the whole "noun chain tree." Such trees have been extensively studied by my colleague Paul Weston, who made a discovery which is of considerable importance in this context. He found that in all Indo-European languages which he studied so far there were never more than seven levels of abstractions that separated the root from its most distant twigs. It is highly suggestive to associate these levels with basic cognitive processes, particularly in view of the fact that the few root nouns that evolve are directly associated with most fundamental cognitive dissimilarities as seen from the earlier examples.




Figure 9. NOUN TREE OF ABSTRACTIONS.

After these preliminaries I am going to launch my last argument, namely that the code of a message cannot be broken unless the internal representation of the common environment is similar—or even isomorphic—in both the transmitter and the receiver. In other words, comprehension is identification with the speaker (7).

Let me demonstrate this now on a minimal model, which again was developed by Paul Weston for demonstration purposes. He, however, studies this mapping problem (ψ in Figure 3) on a large linguistic codex with the aid of a formidable array of computing machinery.

Our problem is simply as follows. Consider a language that has α symbols, and where each message M_i contains n_i such symbols which form m_i words ($n_i \geq m_i$). Let us say that k such messages are received:

$$M_1, M_2, M_3, \dots, M_k.$$

That's all that is known of that language. There are two immediate problems

- 1) What sequence of symbols represents a word in this language?
- 2) What do these words stand for?

In order to follow this argument with greater ease, let us assume that each message contains precisely four symbols ($n_i = 4$, for all i), which form two words ($m_i = 2$ for all i), and there are only two symbols: 0 and 1 ($\alpha = 2$). Assume the message

$$M_1 = "1001"$$

has been received.

Since we know it contains two words we may make several hypotheses about what sequence of symbols may constitute a word. This is quickly done by carrying out all segmentations of the message into two parts ($m = 2$):

$$1001 = \begin{cases} 1,001 \\ 10,01 \\ 100,1 \end{cases}$$

Here the number N_s of segmentations is

$$N_s = 3.$$

However, in the general case we have for a particular message M_i :

$$N_s = \prod_{i=1}^k \left(\frac{n_i - 1}{m_i - 1} \right)$$

$$N_s(i) = \left(\frac{n_i - 1}{m_i - 1} \right),$$

and if all k messages are considered

which soon gets us into pretty large numbers. However, there comes a powerful constraint to our rescue, and that is that every "word" generated by a segmentation is not necessarily a legitimate word. This constraint postulates that any symbol sequence which can be generated by just extending a given sequence cannot be taken as a word. For example, the segmentation of M_1 :

$$1,001$$

represents two legitimate words, for neither word—(1) or (001)—can be generated by extending the other. However, the last segmentation:

$$100,1$$

must be rejected, for the word (100) can be generated by extending (1) by attaching two 0's. The justification for this rule of prohibitive extensions is quite obvious, for any word (S_1S_2) that is an extension of word (S_1) will be decoded as (S_1) and (S_2), and hence leads to ambiguities right from the start.

I may note in passing that the problem of finding in all segmentations the number of those that are excluded by the rule of prohibitive extensions

$$N_s^* = ?$$

is an as yet unsolved combinatorial problem. Of course, everybody here is invited to try his luck. But please drop me a postcard when you have an answer. We would appreciate this immensely.

With these preliminaries we are now ready for a genuine problem.

The old rules hold: "a message has four symbols and consists of two words; there are only two symbols." We have received two messages

$$M_1 = 1001$$

$$M_2 = 0101$$

Question: What are the words?

We immediately proceed in making all permissible segmentations by eliminating all those that violate the rule of prohibitive extension. Figure 10 shows this process in two steps. First the two messages are segmented and the illegitimate ones eliminated(*). In the second step all segmentations are combined and, again, the ones that are mutually inconsistent according to the extension rule are eliminated.

After this weeding operation there remain only three consistent hypotheses as to the vocabulary contained in M_1 and M_2 . One hypothesis, H_1 , says there are only two words, namely:

$$H_1 \dots \dots \dots (10), (01);$$




Figure 10. ALL PERMISSIBLE SEGMENTATIONS OF TWO MESSAGES INTO TWO WORDS BY APPLYING THE RULE OF PRIMITIVE EXTENSIONS.

the other two assume that there are three words buried in these messages:


$$H_2 \dots \dots \dots (1) (001) (01)$$

$$H_3 \dots \dots \dots (1) (001) (010)$$

The question now arises, which school of linguistics is correct? No answer can be given, unless reference is made to the environment in which the partners of this discourse are immersed.

I shall now resolve this puzzle on two levels, both of which do not appear, at least to me, to be trivial.

On the first level, let us assume that there are only two objects, ω_1 and ω_2 , in the universe Ω of the discussants. Then the problem is resolved, in as much as hypothesis H_1 is the only acceptable one, since it postulates the existence of only two words. However, the correspondences



are still in the air.

On the second level, let us assume that there are three objects, ω_1 , ω_2 and ω_3 , in the universe of the discussants. Clearly, hypotheses H_2 and H_3 are now contenders, but no decision can as yet be made as to which is the correct one.

However, no difficulties whatsoever arise if we consider that these two chaps are not just talking. They talk about something relating to the universe of which their cognitive apparatus has produced an internal representation in the form of a structured set of relations Σ_R via the mapping ϕ (see Fig. 3). Let us say that these relations are between the objects, and are of the form:

$$R_1 [\omega_1, \omega_2] ,$$

$$R_2 [\omega_3, \omega_3] .$$

With this the mapping ψ of the messages Σ_M into the internal representations Σ_R becomes a cinch. Clearly

$$R_1 = \psi (M_1)$$

$$R_2 = \psi (M_2) ,$$

with

$$\omega_1 \rightarrow (1); \omega_2 \rightarrow (001); \omega_3 \rightarrow (01) ,$$

and hypothesis H_2 is validated.

With this little *tour de force* I have come to the conclusion of my presentation. I hope that I have mustered sufficient ammunition to defend the position that information storage and retrieval can be looked upon as an exercise in inductive reasoning within the constraints of cognitive processes and linguistic representations. I am, of course, aware of large gaps in my defense, particularly in meshing the theory of relations with the theory of cognitive networks. However, I did not come to say that I have all the answers. I just wanted to respond to Dr. Davis' challenge which, I feel, is legitimate and, moreover, fascinating. The gap, I feel confident, we will eventually close.

Nevertheless, to demonstrate this gap dramatically, let me show you my last picture (Figure 11). This messes up neatly our cognitive processes as well as our linguistic abilities. Fortunately, it does not interfere with our metabolic activities. Hence, I propose, let's join our hosts at lunch which they have so kindly prepared and are now ready to serve.




Figure 11. A DOUBLE PRONGED TRIPLE PRONGED FORK.

REFERENCES

1. J. J. Katz: *The Problem of Induction and its Solution*. University of Chicago Press, Chicago, (1962).
2. L. Wittgenstein: *Tractatus Logico-Philosophicus*. Humanities Press, New York, (1963).
3. W. S. McCulloch and W. Pitts: "A Logical Calculus Immanent in Nervous Activity," *Bull. Math. Biophys.* 5, pp. 115-133, (1943).
4. H. Von Foerster: "Computation in Neural Nets," *Currents in Mod. Biology I*, pp. 47-93, (1967).
5. H. Von Foerster. "Memory without Record" in *The Anatomy of Memory*, D. P. Kimble (ed), Science and Behavior Books, Palo Alto, pp. 388-433, (1965).
6. H. Von Foerster, A. Inselberg, and P. Weston: "Memory and Inductive Inference" in *Bionics 1966*, H. Oestreicher (ed), Gordon and Breach, New York (in press).
7. H. Von Foerster: "Logical Structure of Environment and its Internal Representation" in *In Environment*: Tr. Int. Design Conf. Aspen, R. E. Eckerstrom (ed), H. Miller, Inc., Zeeland, Mich., pp. 27-38, (1963).