NLP 74

OCCAM:
A COMPUTER MODEL FOR A CONTENT ADDRESSABLE
MEMORY IN THE CENTRAL NERVOUS SYSTEM

D. N. SPINELLI

Stanford University
Stanford, California

It is a common observation that man and animals can learn to recognize new situations and objects, a faculty which has fascinated philosophers and scientists through the centuries. In modern times experimentation has supplemented theoretical thinking and the search for mechanisms capable of explaining how sensory patterns are processed and stored in the brain has become a major aspect in physiological and psychological research. The tremendous development of digital computers has influenced theory-making in the biological and psychological fields and has also allowed the possibility of checking theoretical models in actual operation by the simulation of hypothesized neuronal networks or by the use of programs that simply implement the logical operations believed to take place in the process being studied. Whatever the biochemical processes are by which a final mark is laid in the proteic structure of the brain, it is clear that the reading-in and the reading-out of memories has to be effected through the medium of neuronal activity. As the activity of nerve fibers and nerve cells is accompanied by electrical manifestations, it seems at least logically possible that one should be able, so to speak, to surprise memories while they are in the process of being recorded or played back. Efforts in this direction have been numerous; I will therefore refer to only a few that seem to indicate the general findings.

By training animals to differentiate between a 10 per second and a 6 per second flickering light and recording from different brain structures with electrodes, John & Killam (1959) were able to show that the frequency of the stimulus could be recorded in structures which are part of the specific visual pathway and also from structures which are part of the nonspecific system. Of their findings, the most dramatic one has to do with records obtained during

stimulus generalization. The records from the visual cortex of a cat trained to press a bar to avoid a shock whenever a flicker of 4 cps was presented show 4-cps activity. If a flicker of 10 per second is now presented after learning, the animal still performs the avoidance response, but the records from visual cortex now show what look like a mixture of 10 per second and 4-cps activity. After differentiation to the 10 cycles, the activity in visual cortex is mostly represented by 10 cps. It would seem that when the animal is generalizing to the 10 cps, two kinds of activities are generated in the cat's visual cortex: The 10 per second is produced by the stimulus; the other one, the 4 per second, could conceivably be a playback of what the animal had previously learned. Because the animal was expecting 4 per second, he performed appropriately to the 10 per second. The real reason he performed appropriately to the 10 per second was because the previous memory, namely the 4 per second one, was being played back with the appropriate behavior attached to it. Recent work in our laboratory (Spinelli, 1967) also supports the idea that different stimuli generate different forms of electrical activity in the visual cortex of a monkey. Conceivably these different wave forms[1] could be learned by the cortex and played back during recognition. Accompanying wave forms (Pribram, Spinelli, & Kamback, 1967), signaling the presence or absence of reinforcement and the type of behavior that is going to be performed, have also been detected. Again it is conceivable that the total complex of wave forms might be stored by the brain. The representation of stimulus, behavior, and consequences of the behavior would then be available for further reference.

In a different vein, a similar result was also obtained by Morrell (1961) with his experiments on the mirror focus. It has been known for some time (Kopeloff, Barrera, & Kopeloff, 1942; Kopeloff, Chusid, & Kopeloff, 1954, 1955) that an epileptic focus situated in one hemisphere causes, after a period of time, the production of another epileptic focus in the opposite hemisphere at the mirror point. Morrell revived these experiments. Epileptic foci were produced in one hemisphere in rabbits by freezing the cortex. After a week or two, an epileptic focus developed in the opposite hemisphere. If the corpus callosum was sectioned at this time, the secondary focus ceased to exist, but if more time was allowed to pass, about a month or so, sectioning of the corpus callosum did not make the secondary focus disappear. In other words, the secondary focus had assumed a life of its own. The implication seems to be that the healthy tissue in the opposite hemisphere "learned" the pattern of activity that the mirror focus was sending through the collosal fibers. Morrell's result suggested the following experiment: A chronic stimulator was implanted in a cat and a point was stimu-

---

[1] Brain waves are here interpreted to be a more or less direct expression of neuronal processes so that while they may not be direct "carriers" or "codes" of the information transacted they must nevertheless be correlated with the activity of the neuronal networks, much as the noise of a mechanical calculator would be correlated to the operations being performed.

lated at 6 per second in the anterior part of the lateral gyrus. Records were taken from the mirror point. After activating the stimulator, activity in the mirror point was followed over a period of weeks and the stimulator turned off for a few minutes once every few days. Initially activity at 6 per second in the mirror point ceased the moment the stimulator was turned off, but after several weeks of continuous stimulation the 6 per second activity persisted, in bursts, even after the stimulation was discontinued. Again, the indication seems to be that the neural tissue of the cortex can learn a pattern of activity which is repeatedly induced into it and can then play this activity back when appropriately triggered. The question then is: What is the structure of the neuronal network that can so perform? Histological examination of the cortex shows such wealth of connections that it is probably beyond hope to expect to be able to follow fiber after fiber, neuron by neuron, until the whole network consisting of billions of cells is completely unraveled and known. The generation of a parsimonious model would seem therefore not only useful, but indispensable, to allow further experimentation. It is in the nature of a model to generate hypotheses and requirements. The model described in the following paragraphs has highly specific characteristics. Specific types of cells are described and will have to be found to attribute physiological significance to the model. If the elements necessary to the model cannot be identified physiologically, it would then be clear that the model would have to be changed or abandoned. The aim of this model is then to remain as faithful as possible to what is known from neurophysiological studies, but to provide specific assumptions where the data or the theoretical formulation from the physiological field are either incomplete or missing.

NERVE NET STRUCTURE

The speed and reliability with which we recognize spoken words or stimuli presented in any of the sensory modalities make one shy away from a memory model that requires sequential search of any kind among the items stored to identify the stimulus in the outside world. It is therefore assumed, and it would seem indeed desirable, that all memory networks be addressed in parallel by any stimulus entering the central nervous system. In computers an item of information is stored in and retrieved from locations in the core memory: i.e., to retrieve a given item one has to remember where it is stored; this is a memory within a memory requiring indexes and lists. It seems to be more economic to suggest that the basic structure of the memory system used by the brain is not addressed by location (location addressable) but by content (content addressable). What this means is that to retrieve a chunk of information all that is necessary is to provide the system with a fraction of the chunk, and the remainder will be played back. One such network is described here and can be visualized in Fig. 1.
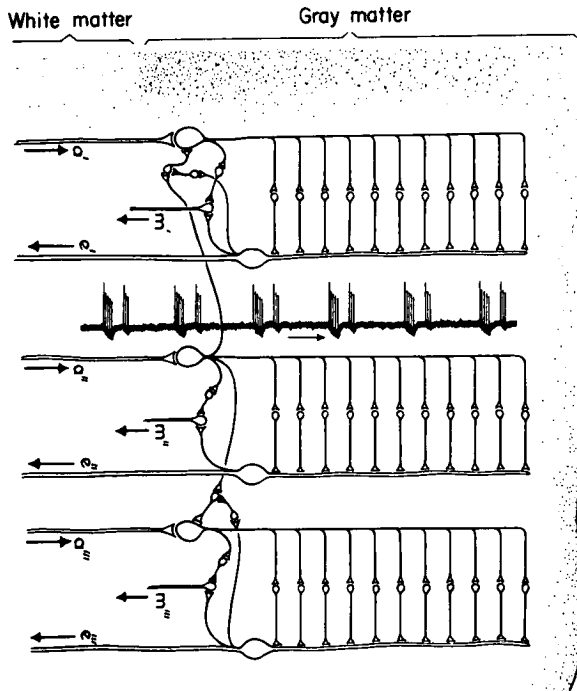
FIG. 1. A network of cells addressed in parallel.

An input fiber enters the cortex and connects itself to a receiving cell; this cell in turn gives origin to many branches that connect themselves to a number of interneurons. The number of interneurons is assumed to be rather great, possibly on the order of several hundreds. The interneurons, in turn, connect themselves to the dendritic ramifications of a further cell. This last cell generates an axon, which leaves the cortex and is therefore part of the output system, whereas the first fiber was part of the input system. Both the input cell and the output cell generate collaterals which connect themselves to a third cell called the Match cell. The Match cell in turn generates an output axon which is also part of the efferent system. This is the basic structure of one content addressable network. In addition, the input cells give rise to collaterals that inhibit laterally other input cells in the nearby networks. Match cells also have collaterals and these collaterals also inhibit input cells in nearby networks.

The Characteristics of Interneurons

It is now assumed that, either because of recovery (cycle-like) or different length in the branches of the input cell or both, the interneurons will be capable

of being activated by the input activity only in a more or less sequential order. In other words, assuming that a pattern of activity is presented to the input cell, this pattern will activate those interneurons that have sufficiently recovered from previous activation, and in turn the interneurons will transmit the activity to the output cell. After a while these interneurons will become refractory or will enter a period of recovery and will be essentially unavailable for further stimulation. Other interneurons will then have sufficiently recovered from previous activity and will transmit subsequent portions of the pattern of stimulation and so on. The logical function performed by the interneurons then is that of a moving window or of a switching network. The final result is that different temporal segments of a pattern of activity will be transmitted by different interneurons. That this switching of activity through different fibers in a very regular temporal manner is achieved in the central nervous system has been shown, for example, in motor nerves, where individual fibers fire at frequencies that are less than 10 per second. Smooth contractions are obtained by regular phasing in and out of the different motor units. A further and crucial assumption to the model is that the synaptic connections formed between the interneurons and the output cell have the characteristic of plasticity. Very simply stated, this means that the more activity is put through one of these synaptic connections, the more open the connection itself will become. Conversely, the less the activity, the less conductive the junction will be. Precisely stated, the assumption is that the synaptic conductivity tends in the limit to be directly proportional to the activity which is going through the synaptic junction itself, so that if a given quantity of activity is presented to the same synapse over and over an asymptote will eventually be reached such that the conductivity will represent faithfully the amount of activity that produced it. The subsynaptic membrane is assumed to have a special characteristic. This characteristic is such that whenever a synaptic connection is activated, the amount of excitatory potential generated is proportional not to activity that generates it, but to the synaptic conductivity. To clarify the function of the interneuron/output cell junction, an analogy is in order.

It is known that muscles become hypertrophic and stronger with exercise, while lack of exercise results in a decrease in the strength and size of muscles. If we consider what is taking place in a single neuromuscular unit, we observe that given a single spike in the nerve fiber the strength of the contraction generated by a single muscle fiber is a function not of the neural spike but of the previous history of the muscle fiber. If the muscular fiber is hypertrophic, a single neural pulse will generate a strong contraction. If the muscular fiber is hypotrophic, a single neural spike generates a small contraction. Exercise, namely the amount of previous activity at the neuromuscular junction, determines the strength of contraction. This is exactly what is assumed to be the function of the synaptic connection between the interneurons and the output cell. To summarize it again, synaptic junctions between interneurons and the output cell have the character-

istic of plasticity. Further, the excitatory postsynaptic potentials generated in
the output cell by activation of one of these junctions is a function of the
previous history of the junction itself rather than of the input activity at that
time. The input activity at that time modifies the junction in the sense that
postsynaptic activity tends, upon repeated presentations, to an asymptote which
represents faithfully the amount of input activity repeated at the same junction.

Match Cell Characteristics

The Match cell receives collaterals both from the input cell and the output
cell and the assumption here is that the Match cell fires only if coincident
activity is received from both terminals. When activity in the input cell and in
the output cell is identical, this cell will fire maximally. When activity in the
input cell and the output cell is completely different, this cell will fire minimally
or not at all.

NETWORK PARAMETERS

Lateral Inhibition and Redundancy of Storage[2]

Let us now assume the existence of a number of these networks. Let us say,
about 50 of them, and let us say that a given wave form, for example, an evoked
potential, is presented to all the networks in parallel to the input fibers over and
over and over. The regular switching of the interneurons assures us that different
portions of the wave form will be stored at different synaptic junctions in all the
networks. Initially all the networks will begin to store the same input pattern,
but eventually one of the networks will be just a little bit better than the
neighboring ones in reproducing the input pattern through its output cell, so
that the Match cell will be more active for this network than for the others. At
this point, lateral inhibition will inhibit nearby networks and will prevent them
from learning this particular pattern any further so that essentially the network
that just by chance gets ahead first will draw the pattern to itself and will
prevent the other networks from learning it. *The number of networks that learn
the same pattern is thus determined by the extent of the lateral inhibition.*
Without lateral inhibition all networks would learn the first pattern presented to
them all. With an infinitely far reaching lateral inhibition, namely with an inhibi-
tion that reaches all the 50 networks, only one network will learn the pattern. If
the lateral inhibition only reaches three or four networks away, five or more
networks might learn the pattern and so on.

[2]The terminology of the following paragraphs has been freely borrowed from neuro-
physiology, psychology, and computer engineering. All terms have retained the correct
meaning.

Afterdischarge and Learning Speed

A second important parameter in the input side of the net is the afterdischarge of the interneurons. It is assumed that each time a cell is activated, an afterdischarge ensues. The longer the afterdischarge, the more the synaptic conductivity will be changed. The shorter the afterdischarge, the less the synaptic conductivity will be modified, so that different times will be needed to reach the asymptote, i.e., different learning rates are possible.

Usage and Novelty

Let us assume that we have presented a waveform a number of times to the 50 networks and that the learning speed, namely the afterdischarge, was such that the waveform has been learned in about 50 or 60 presentations. If the lateral inhibition was sufficiently strong to inhibit the 50 networks on each side of the network that ended up learning the pattern, then only one representation of the waveform will be present in memory. Assume now that a second wave form is presented to the network. If this new wave form is totally different from the one that has already been learned, the Match cells will find no similarity between the input wave and what the net as a whole contains. The content of the net, of course, is of one wave form which we assume to be a meaningful string of signals and 49 sequences of completely random numbers. The chance is then one in fifty that the new pattern will overlay a preexisting memory, and thus destroy it. It seems clearly desirable that *new waveforms*, namely new strings of signals, should be stored *into networks that have not been previously used.*

In essence, it must be that networks that have been used in the past many times are *harder* for a new waveform to enter than are networks that have never been used or have been used fewer times. A simple way of achieving this is to endow the Match cell with some plasticity of its own but of a special kind. We can assume that the Match cell is initially linear in its responding to the number of matches between input and output waveform but that with further usage a nonlinearity begins to appear for low levels of match. In other words, low levels of match would not be signaled by a Match cell that has been activated many times in the past but higher levels of match would be signaled normally. The amount of nonlinearity would have to be proportional to the number of times that the network has been used previously. A new waveform entering the net would then produce in the Match cells a very limited amount of activity which would be due to chance matching between the unrelated content of the memory networks and the values characterizing the waveform. While all previously unused networks would signal chance level of match, previously used networks would signal a below chance level of match. The new waveform would then be stored in that one of the unused networks that by chance presents the highest

D. N. SPINELLI

level of match. This network naturally would inhibit all the other networks and
prevent them from learning the second waveform.

Similarly, a third waveform would again generate a purely chance level of
matches in all the unused networks and a below chance level of matches in the
networks that have stored waveform number 1 and waveform number 2. If
waveform number 1 is now again presented to all the input fibers of the net-
work, it would be immediately recognized; namely the Match cell of the net-
work that stored waveform number 1 originally would present the highest level
of activity of all the Match cells. Presentation of only half of waveform number
1 would still produce a higher level of activity in the Match cell of the network
containing waveform number 1. It can then be seen that such a complex of
content addressable networks is not only capable of pattern recognition, but it is
also capable of playing back the total pattern when it is presented with only a
sufficient fraction of the pattern itself. This last form of behavior can be used in
explaining the associative properties of such a memory network in the temporal
or the spatial domain. It can also explain S-R behavior; namely, if an organism
has had the repeated experience that pressing a red circle is followed by the
appearance of a peanut, whereas pressing the green square is followed by the
appearance of no peanut, then the appearance of the red circle could con-
ceivably generate or cause the playback of the full sequence which involves
appropriate behavior for obtaining the peanut.

Admittance and Generalization

A third parameter is the admittance of the Match cell. This parameter really
defines the amount of variability admitted for each point to be detected by the
Match cell. If the Match cell requires a very low degree of variability between
each point of the input wave and *each point* of the output wave to be activated,
then we could say that the admittance is very, very small. Conversely if the
degree of variability allowed by the Match cell between the input and output is
great, we could say that the degree of admittance is greater. In other words, this
parameter has to do with the y dimension of the wave and determines the range
over which y values will be *generalized,* i.e., considered to be the same.

Acceptable Match and Risk

The fourth parameter controlling the net is the acceptable match. This para-
meter, not to be confused with admittance, has to do with how many points
between the input and the output waves were found to be matching by the
Match cell, i.e., the x dimension of a waveform. If all the points were matching,
then the match is, of course, 100%. If only half of them were matching, then the
match is 50% and so on. It is clear that the total presentation of a pattern which
has been learned before will provide 100% match and therefore full recognition.

But this condition is also the least informative. In other words, while there is no uncertainty about the pattern, there is also no extra information furnished by the recall of this pattern. A less than total presentation will provide the organism with some extra information, namely the stimulus from the environment will generate the playback of a sequence of signals which is longer than the sequence provided by the stimulus itself.

If wave forms in the brain represent stimuli, responses, and the consequences of responses as we have previously seen (Pribram *et al.*, 1967), then presentation of the stimulus will generate a playback of the whole sequence; that is to say: recognition of the stimulus, the appropriate behavior that went with the stimulus, followed by the expectation of the consequences of the behavior. The amount of extra information obtained by the network or by the organism is greater, the smaller the segment of the total input string. The amount of uncertainty, and therefore of risk for the organism in using the sequence itself becomes, on the other hand, correspondingly greater. An analogy in the auditory mode helps in understanding the significance of this parameter. The name of a song followed by the playing of the whole song will, of course, be recognized, if it has been heard before. The name of the song followed by half of the song will enable the listener to remember the remainder of the song. Ultimately, just the name of the song, or a few notes, will enable the listener to recall it entirely. But if the notes are too few, or if the name of the song is equivocal, then the level of match would be correspondingly very, very small and might not enable the recaller to identify which song we are referring to. It might be that the few notes provided are part of the beginning of many songs. Ideally then, the acceptable match parameter should be set for that minimum value which allows unequivocal recognition of the stimulus with recall of the associate behavior and consequences of behavior. It follows from the above paragraph that such a network is then capable not only of pattern recognition and of S-R behavior, it is also capable of being biased. Imagine a number of strings of signals recorded in memory having to do, for example, with feeding behavior. These strings would be located in visual memory for those parts of the feeding behavior that are directly connected to vision, for example, seeing an appropriate stimulus that when manipulated under visual control leads to availability of food which then can be taken to the mouth, be ingested, and produce subsequently the cessation of hunger. The assumption has to be that while visual memory contains most precisely and primarily visual strings, it also contains enough nonvisual information to allow the readdressing of the system by the visually triggered strings so that auditory, somatic, gustatory, etc., strings are subsequently called into play. The internal state which would be part of the string, for example, hunger and the disappearance of hunger, would activate or would facilitate all those memory strings that contain such information in themselves and therefore produce a partial level of match. This would then make available to the rest of the brain strings containing pertinent information about feeding behavior. If other parts of

some strings are available in the environment, a higher level of match would be achieved for certain strings and the connected behavior could then be played back if the acceptable level of match is reached or exceeded.

An analogy at this point again might help in understanding the functions of a network containing, for example, 50 content addressable units. Imagine that we have 50 computers for average transients (CAT) and let us say that at zero time their memory cores contain nothing but random numbers. If we now present to the 50 CATS in parallel an input wave over and over, an average of the input wave will begin to be formed in all 50 of them. The one CAT computer that begins to show, ahead of the others, more points of match between the input wave and the output, can inhibit the remaining 49 and will prevent them from proceeding with the average, whereas it will continue to build up a more and more clear representation of the input in its own core memory. The decision on which CAT computer is ahead of the others is performed by the Match cells. Match cells, in other words, determine the degree of match between input and output for each memory network. The CAT computer analogy is also helpful in understanding that a small or medium amount of variability in the input waves will not upset the network but will still result in the learning of a waveform which is the representative average of the waveform presented.

COMPUTER IMPLEMENTATION OF AN OMNIUM-GATHERUM CORE CONTENT ADDRESSABLE MEMORY (OCCAM)

Simulation of the network described above was undertaken to verify that the hypotheses and assumptions made would indeed enable the network to perform the functions required. The simulation was performed on a small general purpose computer, the PDP-8. The core memory of this computer consists of 4096 twelve-bit words and is therefore too small to allow simulation of *all cell characteristics*. Only those functions of each cell that are involved specifically in the model have therefore been simulated. With respect to interneuron function, it is clear for example that the presence of the input cell is not really necessary and that a fiber branching onto the interneurons would do just as well, provided one is willing to accept presynaptic inhibition rather than postsynaptic into the system of lateral inhibition. Similarly the interneurons themselves are not really necessary, provided one is willing to assume that the ramifications of the input fiber, because of different length and diameters, act in effect as a switching network. The fundamental characteristic is the presence of lateral inhibition. Without it the network would only learn one pattern of activity regardless of the number of content addressable networks contained in the total net. As has already been said, the extent of lateral inhibition determines the amount of redundancy with which a pattern is stored in the net.

The second set of crucial elements in the neuronal model are the synaptic connections between the interneurons and the output cell, or alternatively between the ramifications of the input fiber and the output cell. It is this junction with its special plasticity and subsynaptic membrane characteristics that together with the characteristics of the Match cell make it possible for the model to perform.

What has therefore been simulated is the following: a switching network which breaks an input wave into fifty segments. Each segment then is directed to a "synaptic" junction. The special plastic characteristics of the junction are simulated. The postsynaptic junctional activity is adjusted each time by a small fraction so that if the input to it is repeated a number of times, it would eventually asymptote to it. The activity generated by the postsynaptic membrane in turn depends only on the past history. In other words, the average of all preceding activities at the synapse is generated every time the synapse is activated, irrespective of the amount of activity causing the activation.

Concerning the Match cell, the characteristics that have been simulated are the admittance, namely, the amount of similarity between each one of the 50 input values and the 50 output values that is required. In addition, a usage factor which determines output on the basis of the number of times the output cell was activated. Finally, the property of totalizing the number of matches detected has also to be simulated for this cell.

Of the parameters described for the neuronal model, four have to be given real values, for example: lateral inhibition, i.e., redundancy of storage, had to be set at one value or another and in the actual simulation it was set for a value of 50 which means that only one representation of each wave form would be formed. The afterdischarge and therefore the learning speed (namely, the number of times the same pattern afterdischarges within the networks) were set in such a way that learning would take place in 50 to 60 trials.

A Match cell usage parameter, as already noted, is the inverse measure of novelty. Admittance value, i.e., generalization, also has to be specified together with the parameter that controls the increasing nonlinearity. These two parameters together control the "tuning" with which a given network recognizes its own wave form. The fact that the networks which have been used many times are harder to enter gives this program not only the characteristics of a content addressable memory but also the characteristics of a push down list. Assuming that memory were ever to be filled, a new wave would be stored in the network that has been used the least of them all, which makes sense from an adaptive point of view because the least used string is, by definition, the least useful. The admittance parameter is crucial in determining the way in which a waveform is learned. If the parameter is too loosely set, if the admittance is very broad, then subsequent waveforms will overlay previously learned ones, because the tuning of the networks is not fine enough. At the other end of the spectrum, if the admittance parameter is too narrow, a given network will not be able to recog-

nize on the second presentation of a waveform that it is the same waveform that it began storing on presentation number 1. Therefore a new network might take on the waveform at the second presentation, but because a very small amount of learning takes place from one presentation to the next, this second network might not recognize the waveform on the third presentation, so that a new network would come into play for each presentation. In other words, the net would be prevented from learning the waveform because it could not, by using a too rigid criterion, recognize it until it was fully learned, which would effectively make it incapable of learning.

One way out of this impasse would be to maintain a very narrow acceptance but increase the learning speed so that instead of 50 to 60 presentations, only two or three would be required. In this case after one presentation, learning would have proceeded so fast that there would be no question about recognizing the waveform on presentation number 2. This solution is far from ideal for a number of reasons; the first one being that slow learning imparts to the system *averaging* properties which seem to be a desirable feature. The second reason is somewhat connected to the averaging properties and would be that slightly different versions of the same waveform would be stored as two separate representations, namely the system would be unable to generate representative strings of signals rather than individual ones and therefore would store many representations of the same pattern.

The fifth parameter, namely, acceptable match, does not need a real value and can be set to zero. When this is done, Occam simply shows the best match it finds in its networks to whatever waveform is being presented to it. It is clear that in a real organism, the value of the acceptable match would therefore have to be defined by the software in some fashion.

Figure 2 shows this program, named Occam for Omnium-Gatherum Core Content Addressable Memory, in real operation. A waveform in A is presented to Occam repeatedly. From the top are subsequent responses, every 10 trials showing further and further improvement until after about 50 trials Occam generates a waveform which is practically identical to the input wave form. In Fig. 2b, Occam is presented with a second waveform which is also learned. Upon presentation of parts of waveform number 1 or waveform number 2, Occam retrieves the remainder of the waveform appropriately.

GENERAL COMMENTS

Reinforcers as Dissimilarizers of Memory Strings

It is evident from the model that Occam would have a hard time learning two different patterns which are very similar to each other. As a matter of fact, this
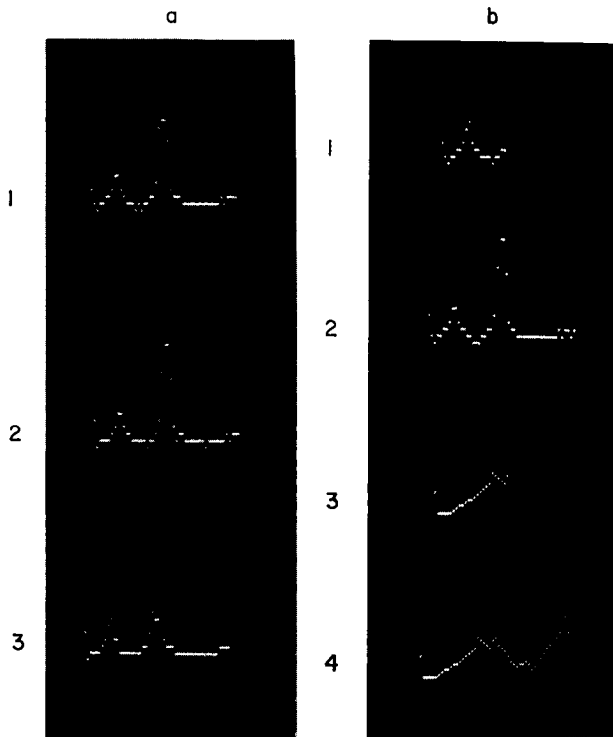
FIG. 2. An Occam program.

is an impossible task. This is because, if two patterns are very similar, presentation of one of them will produce a high level of match in the network that has stored the other one, so that the second pattern would always enter the network of the first one and interfere with it. A way to enable Occam to learn two very similar patterns would be to attach to one of the patterns a different ending signal. For example, we could follow one of the patterns with a series of oscillations at a given frequency whereas we could follow the other pattern with a series of zero levels or with a series of oscillations at a frequency different than the first one. The operation seems to be extremely similar to the one which is done usually in behavioral experiments when an animal is trained to discriminate between two patterns. One pattern is usually followed by a reinforcer whereas the other pattern is not. This serves as a pointer to notify the organism that two patterns which might have looked identical are really different, even though the difference might be small. Similar concepts have been expressed by Pribram (1963). The organism can then institute such procedures that it can look for differences and end up by storing the two patterns or enlarged versions of small sections of the two patterns.

Reinforcers as Controllers of Learning Speed and Redundancy

It is also possible to imagine that reinforcers act on the parameter of learning speed and lateral inhibition. If reinforcers would possibly decrease the extent of lateral inhibition and increase learning speed, this would enable an organism to learn faster and more redundantly strings of signals which are associated with information which is of survival value. It would seem at least at first glance that reinforcers are divisible into two classes. The first kind of reinforcer would be connected to pain, food, and the like. This particular system can be conceived as permanently wired-in so that whenever activated lateral inhibition and learning speed are appropriately affected. The second kind of reinforcer, having to do with social situations, psychological situations, and the like, would be acting on the memory only through the software, namely, as parts of existing programs or plans (Pribram, 1963).

To summarize, a computer simulation of a hypothetical neuronal network is provided. The network consists of many identical subunits which are all addressed in parallel and have the characteristic of content addressability. It would seem that the model furnished can explain pattern recognition and stimulus-response behavior.

REFERENCES

John, E. R., & Killam, K. F. Electrophysiological correlates of avoidance conditioning in the cat. *Journal of Pharmacology and Experimental Therapy*, 1959, **125**, 252-274.

Kopeloff, L. M., Barrera, S. E., & Kopeloff, N. Recurrent convulsive seizures in animals produced by immunologic and chemical means. *American Journal of Psychiatry*, 1942, **98**, 881.

Kopeloff, L. M., Chusid, J. G., & Kopeloff, N. Chronic experimental epilepsy in *Macaca mulatta*. *Neurology*, 1954, **4**, 218.

Kopeloff, L. M., Chusid, J. G., & Kopeloff, N. *A.M.A. Archives of Neurology and Psychiatry*, 1955, **74**, 523.

Morrell, F. Lasting changes in synaptic organization produced by continuous neuronal bombardment. In J. F. Delafresnaye, A. Fessard, & J. Konorski (Eds.), *Symposium on brain mechanisms and learning*. Oxford: Blackwell, 1961. Pp. 375-392.

Pribram, K. H. Reinforcement revisited: A structural view. In M. Jones (Ed.), *Nebraska symposium on motivation*. Lincoln: University of Nebraska Press, 1963. Pp. 113-159.

Pribram, K. H., Spinelli, D. N., & Kamback, M. C. Electrocortical correlates of stimulus response and reinforcement. *Science*, 1967, **157**, 94-96.

Spinelli, D. N. Evoked responses to visual patterns in area 17 of the rhesus monkey. *Brain Research*, 1967, **5**, 511-514.