IC/72/83

INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS

INFORMAL MEETING ON NEURAL NETWORKS

24 - 26 July 1972

(SUMMARIES)

INTERNATIONAL ATOMIC ENERGY AGENCY

UNITED NATIONS EDUCATIONAL, SCIENTIFIC AND CULTURAL ORGANIZATION

1972 MIRAMARE-TRIESTE

IC/72/83

International Atomic Energy Agency
and
United Nations Educational Scientific and Cultural Organization

INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS

I N F O R M A L   M E E T I N G

O N

N E U R A L   N E T W O R K S

24 - 26 July 1972

(SUMMARIES)

MIRAMARE - TRIESTE
August 1972

Please note that copies of papers referred to may be obtained direct from the authors and not from the ICTP.

NON-LOCAL MEMORIES IN PHYSICS AND BIOLOGY

C.J.H. Watson

Merton College, Oxford, UK.

There is a sense in which any system which is capable of memorizing information necessarily does so by localizing it in space. Only objects which preserve a certain amount of spatial structure are capable of recording information in a manner which endures in time. However, most of the objects which have traditionally been used for this purpose - objects such as books, photographs, tape recorders and digital computers - store the information in a manner which is "local" in a more radical sense than this: they depend upon the idea of mapping the information, regarded as a sequence of distinct "items", onto a sequence of distinct "locations" within the object, and the spatial relationships between these locations remain constant in time and play a significant role in the manner in which items of information are assigned to locations and in which they are subsequently retrieved. In a computer memory, for example, each bit of information is assigned to a magnetic core which has a constant and significant location within a board of electronic circuitry. However, in principle, it is possible to use objects which have very little enduring spatial structure for information storage , or even to store information in such a way that no one-to-one correspondence can be set up between items of information and locations within the memory, and such systems are "non-local" in the sense of the present paper.

A simple example of a storage device which is "non-local" in the present sense can be derived from the idea of coding. Suppose that the basic "item" of information is a word, and that the body of information to be stored consists of a string of words. Consider a "coding device" which takes the letters of the string of words and permutes them in accordance with some complicated preassigned and reversible rule, and then records the resulting encoded string of letters on a piece of paper. That piece of paper preserves a perfect "memory" of the information, yet within its record there is no definite location at which any word has been stored: rather, each word has been spread over several locations. Thus, this "coding device", regarded as a store for words, is non-local even though it is a local store for letters. Another, slightly different, type of non-local store is one in which successive words are assigned to sequentially numbered boxes, which can then be stacked and restacked in a random manner. Examples such as these show that non-local storage is possible: they do not show that it has any advantage over local storage - indeed, they suggest the contrary, since it appears that the information is only usable after it has been decoded and transferred to a local store.

The potential advantage of non-local memories only appears if one considers the procedures which might be used to retrieve a selected part of the stored information. In local memories, this is most naturally done by searching in the location where the item of information is known to have been put. This procedure, known as "location-addressing", is only possible if quite elaborate arrangements are made at the time that the information is stored - arrangements which involve entities such as indices and files. Local memories do not have to be used in this way: instead, one can assign information to locations randomly, and then retrieve it by searching the entire memory for an item which matches some desired item in its "content", but this is very inefficient unless very powerful parallel search facilities are available. The advantage of a non-local memory is that in principle it is possible to perform such "content adressing" of the information directly, without even localising the information prior to retrieval.

The simplest example of a non-local content-adressable memory is a hologram. From the present point of view, a hologram consists of the interference pattern created by light from two objects which represent information. Mathematically, the record consists of the Fourier transform of the spatial correlation function of the two information patterns. The record is non-local in the present sense - there is no one-to-one correlation between items of information in either pattern and locations on the hologram. It is "content-adressable" in the sense that when it is illuminated by light from one of the patterns, it produces an image of the other pattern.

In 1968 Longuet-Higgins proposed a device which he termed a "holophone", which would record temporal information in exactly the same way that holograms record spatial information. The device records the Fourier transform of the autocorrelation function of the temporal information sequence: when supplied with a "cue" consisting of some part of that sequence, it responds by playing back the remainder of the sequence,

starting from the point "cued". This procedure clearly makes the temporal information centent adressable. In 1970 the author showed that a rather wide class of physical systems possessing some non-linear element and a certain natural periodicity could exhibit holophonic behaviour. The essential feature turned out to be that the system should respond to a signal consisting of three pulses at times t_1 , t_2 and t_3 by producing an "echo" pulse at $t = t_3 + t_2 - t_1$, with an amplitude proportional to the product of the amplitudes at the times t_1, t_2 and t_3 . It is readily shown that such a system produces a playback, when operated in the manner described by Longuet-Higgins, which is mathematically identical to the expression which he derived. The proof turns on the fact that after the recording phase, and before the cue, the system preserves a record of the autocorrelation function $A(\tau) = \int R(t)\, R(t + \tau)\, dt$ of the recorded signal $R(t)$. The playback $P(t)$ in response to the cue $C(t)$ is then found to be proportional to the convolution of $C(t)$ with the auto-correlation function

$$P(t) = \int C(t - \tau)\, A(\tau)\, d\tau$$

which is in turn proportional to $R(t - T)$, where T is the (quite arbitrary) time interval between the cue and the passage cued.

Unfortunately, as Longuet-Higgins and others have shown, the signal/noise ratio of the playback from such holophones is very poor unless the length of the cue is comparable to the length of the entire body of information stored. The reason for this can readily be seen if one uses the "echo" interfunction to predict the amplitudes of the holophonic and noise playback signals. In Fig.1, the N-pulses of a record sequence are indicated, followed by J cue pulses, and a simple diagram is introduced to interpret the echo time ratio $t - t_2 = t_3 - t_1$ - the lengths of the two "bridges" have to be equal.

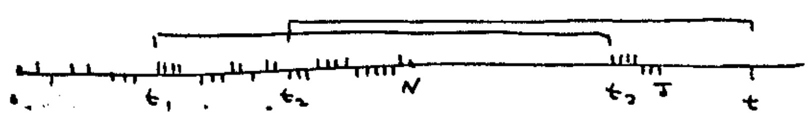

For each of the J positions of the left-hand bridge, the amplitudes at t_1 and t_3 are the same, so the echo pulses at t have the same as, and are proportional to, the amplitude at t_2. These J pulses add arithmetically, and this is true for every position of t_2 within the record. Thus, the holophonic playback at every time t after the end of the cue is proportional to the amplitude of the record following the end of the cued passage. However, for every such "holophonic" diagram there are roughly N diagrams in which t_1 lies at some arbitrary point within the record and t_2 is chosen so that $t - t_2 = t_3 - t_1$, so even if the record sequence averages to zero, the signal/noise ratio is only $J/\sqrt{NJ}$, which is very poor.

It will be seen that the reason for this poor performance is that the left-hand bridge, which runs along acting as a "correlation bridge", in effect constructs the binary cross-correlation of the cue with each passage of equal length in the record. This is rather an insensitive measure of correlation and, for example, only decreases by one part in J if one bit fails to match. By contrast, if the echo rule could be represented by a multiple bridge, such as that illustrated in Fig.2, with M correlation bridges of equal length:

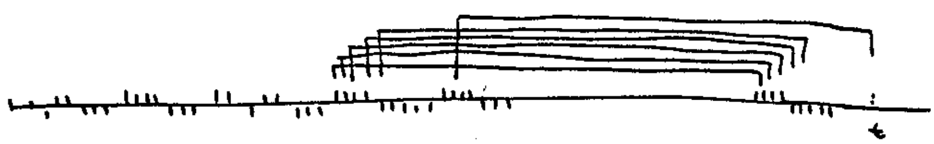

Then if $M = J/2$, each false match between a cue bit and a record bit would halve the holophonic response. By a similar argument to that used in the three-pulse case, it can be shown that any device which stored a "high-order" autocorrelation function such as

$$A_M(\tau) = \int R(t_1)\, R(t_2) \ldots R(t_M)\, \delta(t - t_1 + t_2 - t_3 - \ldots + t_M)\, dt_1\, dt_2 \ldots dt_M$$

would respond in this way. The calculation of the signal/noise ratio in this case is a non-trivial combinatorial problem, but it can be shown that in some circumstances it is favourable. It could be made more favourable still if the echo rule prohibited echoes in which the spacing of the pulses within the record was much larger than that within the cue.

A physical system has been found - a collisionless plasma - which can function as a high-order holophone of this kind, albeit rather imperfectly. It is tempting to speculate whether neural nets might be able to operate in this way. Two rather different possibilities have to be considered: either the net could store directly, in a manner somewhat analogous to a tape recorder, the high-order autocorrelation function $A_M(\tau)$ (or some more suitable generalization of it), or it could behave like all existing holophone implementations and store the Fourier transform. In the latter case it would be necessary to find some non-linear elements within the brain which possessed the rather strict periodicity at a high frequency which this requires. It seems conceivable that such periodicities might be found in biochemical cycles within a cell, and that this might be a mechanism which related to the long-term memory of the brain.

REFERENCES

H.C. Longuet-Higgins, Nature 217, 104 (1968).

C.J.H. Watson, Nature 229, 28-30 (1971); "The plasma echo holophone", Culham Lab. Report CLM-P317, 1972 (unpublished).