

10P
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Artificial Intelligence
Memo. No. 48

March 1963

Neural Nets and Theories of Memory

Marvin Minsky

*This empty page was substituted for a
blank page in the original document.*

MINSKY: *Neural Models for Memory*.

A number of models developed in work often called "neural-net" research may be of interest to physiologists working on the problem of memory. From this work comes a variety of ideas on how networks of neuron-like elements can be made to act as learning machines. Some of these may suggest ways in which memory may be stored in nervous systems. It is important, perhaps, to recognize that these models were not founded at all on physiological ideas; they really stem from psychological and introspective notions. They all involve some form of alteration of synaptic transmission properties contingent on the pre- and post-synaptic activity during and after the relevant behavior. This notion is suggested not so much by actual observation of synapses as by the introspective simile of wearing down a path—the "ingraining" of a frequently-travelled route. Below we shall argue that this idea is useful and suggestive, but not sufficient. These models can be made to account for learning connections between stimuli and responses on a low level, but do not seem to account for higher, symbolic, behavior. We will argue that the latter suggests a return to the search for localization of memory, a topic that has been unpopular for many years.

I. *Early neural-network models*

It would take a good deal of space to discuss all of this work; we can give only an outline of some of the major steps. This discussion is not intended to be a thorough review, and we discuss only those models connected with theories concerning memory. That is why there is no reference to other theories of model neurons, e.g., those of Harmon. A variety of mechanisms were proposed in the 1930's by Rashevsky and his colleagues (see, e.g., ¹ and ²); these models were based on a threshold neuron with excitatory and inhibitory inputs that summate with exponential decays. In 1943 McCulloch and Pitts published their analysis of the logical domain of some even simpler "neurons"; the stimulation for these has a simple all-or-none character and the cells have a simple threshold with absolute inhibition³. Nevertheless, we see that in principle, at least, networks of even such simple cells can exhibit adequately complex behavior, when properly assembled. This analysis was completed by Kleene⁴ and the clearest account of the results will be found in Copi, Elgot and Wright⁵.

The networks constructed in this area are very highly constrained and sensitive to wiring errors and alterations, suggesting that this kind of model is incompatible with a physiological theory. To correct this, von Neumann⁶ developed networks whose input-output characteristics were insensitive to random *independent* fluctuations of the cells. More recently, McCulloch and his colleagues⁷ have shown that one can make some such networks insensitive even to *non-independent* fluctuations of the cell properties. See Cowan (this volume).

The networks of 6 and 7 are insensitive to some fluctuation of cell properties at a microscopic level, but the behavior of such networks is still very precisely dependent on the details of the interconnections. This is not consonant with our current picture of the structure of the brain. While we are finding each year more and more order, there still appears to remain a great deal of connection unspecificity. This raises the problem of how one can obtain orderly behavior and learning from an initially weakly-constrained structure. This problem led to the construction of a series of "random-net" models which begin with a network of elements arranged in an initially unknown, rather disorderly structure. Obviously, one would not propose this as a complete brain model, but it is a good medium for studying this aspect of the problem.

One of the earliest, and still the most ambitious and elaborate of the random-net models is that of Hebb⁸. In this theory, pathways are selected and facilitated as a result of certain activity patterns, and this leads to the formation of certain more-or-less circular reverberatory patterns called "cell-assemblies". This proposal had its antecedents, but was nowhere else developed to the extent described in ⁸. We will return to this model later.

A learning model must account for simple forms of reinforcement learning. To do this one must have means for generating a variety of reactions and a scheme for selectively emphasizing those correlated with successful or rewarded behavior. Probably the earliest random-net system in which this could be demonstrated is that described in Chap. IV of Minsky⁹. In this machine behavioral variation is generated by assigning a transmission probability to each synapse. The effect of reinforcement is to modify the transmission probability of those synapses which have recently succeeded in exciting the post-synaptic cell. This has the effect of a probabilistic selection of stimulus elements, along the lines of the theories developed by Estes. (As I went to so much effort to allow for a wide range of probabilities, I am singularly intrigued by the paper Estes has just presented; it appears to show that a very much simpler structure might yield equivalent results.) This probabilistic random-net machine, called the SNARC, was able to learn some fairly complex discriminations, and to find its way through quite complicated mazes (when given different stimulus patterns for the different vertices). As it was able to establish circular internal pathways it could also learn some limited sequential discriminations. Nevertheless, the experiment convinced me that the real problems lay not in the source of variation but rather in the mechanisms for assembling hierarchies of behavioral elements. For this I found it necessary to turn toward models more like those of Hebb. This later work reported in the latter chapters of ⁹, led to some schemes that might obtain (from some not-so-random nets) various forms of prediction, expectation, and planning. A discussion I have just had with Anokhin suggests that he had reached related conclusions quite a long time before this.

The use of physical hardware makes research in this area expensive and inflexible. The first reported experiments on learning in random nets, using a digital computer, are those of Farley and Clark^{10, 11}. Here, behavioral variation is introduced through fluctuating thresholds. A Farley-Clark cell is somewhat like a Rashevsky cell; it fires when the excitation exceeds a threshold. The signals pass through synapses that attenuate the strength of the signal, and learning is mediated by modifying the attenuation coefficient for each synapse participating in the pre-reinforcement reaction. Again, because the random net allows for circular reverberation, it can learn some discrimination of temporal patterns. Following earlier work of Beurle (see ¹²), Farley¹³ has also studied the problem of activity patterns in large randomly connected networks; while there is no learning in these experiments, it will certainly be necessary to understand these results if one is to discover how to make stable large random networks. In this connection one should know also the related paper of Selfridge¹⁴.

The next development in computer simulation of random nets is reported by Rochester *et al.*¹⁵ who describe attempts to simulate cell-assemblies. The result is that one can find connection constraints and cell properties that do lead to assembly-formation in nets that are still quite randomly connected. The results in ¹⁵ do not go far enough to suggest that the assemblies themselves can become hierarchically interconnected as suggested by Hebb. Further theory along these lines, but without experimental confirmation is reported by Milner¹⁶.

In the last few years, there has been much activity concerned with the study of certain much simpler networks. These have random connections from one layer of cells to another, but no return connections. The first of these is the "Perceptron" model of Rosenblatt^{17, 18}. The synapses are like those of Farley and Clark¹⁰ with reinforcement modifying those attenuation coefficients, or connection-weights, which participate in each rewarded decision. Alternatively, one may use negative reinforcement only for error-correlation. The network is set up to select one of the cells of the output layer, along the lines suggested by Selfridge¹⁹ to represent discrimination of the stimulus as one of a number of categories. That output cell which receives the largest excitation dominates the others; in some models through a cross-inhibition scheme, in others by a retroactive inhibition of the first layer. These nets can be made to learn certain discriminations, but unless the network is preceded by a sophisticated stage of pre-processing (e.g., like those suggested by Hubel and Lettvin *et al.* on the physiological side, or Von Foerster and his associates on the synthetic side) they cannot learn to make generalizations beyond those entailed by the overlapping of similar stimuli (as discussed in Clark and Farley¹¹). This seems to be a fundamental limitation; recognition of generalizations within this domain is quite valuable, but to obtain sophisticated symbolic behavior one must go beyond it. In these machines one can interpret memory as taking the form of storage of empirically-estimated conditional probabilities; an analysis of this (and an evaluation of this family of neural-net models) is found in Minsky and Selfridge²⁰. In that connection one should consult also the work of Uttley²¹, although that is not a *random* net approach.

Closely related models are found in the family of elegant devices of Gamba²²; here we find decision processes based on correlations with randomly generated templates, with learned conditional probability weightings.

All these devices appear to have considerable capacity to learn to discriminate between sets of stimuli which have actually been presented. They appear to be much less successful at generalizing the discrimination to stimuli which have not been previously presented. And contrary to published claims, there is no evidence that this limitation is automatically transcended by going over to very large networks of the same kind. In this regard, see the recent critical paper of Bryan²³.

II. *Current work in machine learning and problem-solving*

We observe that in neural-net research, over the past decade, there has been a trend toward a lower level of aspiration. At the beginning, there was a distinct hope that a very large, highly-connected, network could organize itself to perform sophisticated cognitive activities. One finds today that most work is directed toward obtaining relatively simple discriminations in nets with one layer of connections. This happened because the complex nets could not be made to work without further constraints, and suitable constraints were not sufficiently well understood.

In the same era, we note a closely related field where progress has become remarkably rapid. This is the domain of machine learning and problem-solving. (See, e.g., Minsky^{24, 25}.) As this will be the subject of a paper shortly to be presented by Newell, I will not discuss it in detail, but I want to point out a curious contrast between this and the

neural-net areas. While the work on neural-net models was, in effect, backing away from the hierarchies of symbolic representations proposed by Hebb, the work on computer programs for learning and problem-solving was moving rapidly toward just such hierarchies. Today we have computer programs which do solve problems of considerable intellectual difficulty, and even in those instances where the research effort was directed towards "artificial intelligence"—that is, towards making machines solve intellectual problems without any attempted constraint to simulate human thought processes—we often find the behavior to be strikingly suggestive of that involved in thought. The behavior is highly dependent on the mode used for internal symbolic representation—corresponding to the use of language in reasoning, and on the methods used for storage of partial results and postponement of sub-problems. Now the thing that concerns us is this: in all the really successful experiments in this area, we find certain common features concerning symbol-manipulation. These are discussed in Newell's paper (this volume) and I agree quite completely with his conclusions. Now, what do these conclusions suggest about the brain? There is certainly no *logical* implication at all, for the computer results are based not at all on physiology, and only remotely based on psychology. It is conceivable that the brain works on some utterly different basis. But while this may be conceivable, we simply do not have any such alternative before us today. As Newell has said, this is "the only set of ideas that exists today about how to build these very complicated structures". Since symbols appear to be so important and necessary, it seems compelling that we at least consider experiments to find how they might be represented as brain events.

The same "set of ideas" suggests that memories themselves must be represented as symbols, or as symbolic expressions. The "symbol", as it occurs in our computer programs, is a relatively concrete thing; it has a "location in memory". Now this is not logically necessary; it is "conceivable" that it could be represented only as an *emergent*—entailed by the joint activity of things associated with it, or that it could be represented by some global form of activity, such as a wave interference pattern. The trouble with this is that it is today a useless hypothesis—we have no associations with it and cannot use it to promote further thinking or design experiments. Therefore, distasteful as it may seem to physiologists, the current situation suggests a renewed effort to find memories deposited in something resembling a spatially localized form.

III. *Localization of memory*

It might be reasonable to look again for localization of memory in the brain. I know that this idea is unpopular today; half a century of efforts to locate memories have failed. But it may be useful to review the nature of this failure in the light of our current ideas about the representation of mental events. Fifty years ago, when the neuron doctrine had just taken hold and modern experimental psychology was in its early stages, it seemed reasonable to look for the sites of the changes associated with, e.g., the formation of conditioned reflexes. Designing experiments along this line met with insuperable difficulties, and even in the more modest search for regional localization of broad functions the results were, generally, equivocal. By the 1940's there was a general feeling that such attempts were futile. This attitude, crystallized in the person of Lashley (see, e.g., ²⁰), gave rise to the radical view that it was perhaps hopeless to look for specific loci of memories and other units of intellectual activity, and suggested that the nervous system functioned somehow through the interaction of superimposed gross modes of activity, perhaps wave-like interference patterns. At the time there was no real evidence for such a picture, nor any theory of how such patterns could be organized into an intellectual hierarchy of function but, *faute de mieux*, this view became popular. Today we could, perhaps, construct a better argument along these lines*, but our new ideas suggest equally a re-examination of the localization idea.

* E.g., using the notions of Cowan (this volume).

The negative evidence for localization can be divided into two families; that concerned with the result of *ablation* and that concerned with *stimulation*.

a. Ablation problems and the requisite redundancy

With certain important exceptions, notably in connection with speech, it proved impossible to elicit highly specific memory defects by ablating small, or even large, areas of brain. One found either a gross interference with a function and consequent general deterioration, or else no measurable deficit whatever. The suggestion seemed inevitable that each memory was distributed, more or less uniformly, over the whole brain.

A more moderate, intermediate possibility seems to have been overlooked; it is important because the totally distributed model is probably unworkable. If each memory record is stored in one brain site, and we remove half the tissue, we would expect to remove half of the records. If each memory were recorded in *two* separate sites and we remove half the brain, then we can expect to delete *one-fourth* of the records. If each record is stored in *three* sites, then removal of half the brain will, on the average, remove only one-eighth of the records. This idea of redundant storage is very widely known, yet it is still very poorly appreciated. It makes it seem unnecessary and extravagant to go over to a theory in which the records are copied in an infinitely-partitioned distributed fashion. Thus, suppose finally that each record is stored in no more than 10 places. Now if we remove half the brain, the probability of totally ablating any particular record is less than $1/10$ of 1%! That is, the odds are less than one in a thousand that destruction of half the structure will get all ten representations of any particular record!

None of our clinical tests are sensitive enough to demonstrate an intellectual deficit of less than a few parts per hundred, so that a redundancy factor of so little as 5 or 6 would probably account for those results of moderately extensive brain injury in which no deficit is apparent. It is not my intent to propose that each is in fact really stored in some very narrowly delineated spatial site, with exact copies in several other locations, but only that we ought to turn our thinking back in that direction. In neurophysiology, we have not yet come to appreciate the full force of small amounts of redundancy, though this probably is one of the outstanding properties of the machinery with which we deal. Most of us have not realized the amazing power of a redundancy factor of 10. Factors of 5 or 6 would regularly defeat our deficit tests whenever the patient retains his ability to make common-sense deductions from related data—"well-founded confabulation", one might say.

b. Stimulation and the problem of adequate excitation

A second reason to doubt the localization of memory is the difficulty encountered in trying to elicit memories through direct stimulation. We can probably discount as exceptional those striking hallucinatory incidents occasionally elicited; these probably represent some different mechanism not dependent on specific excitation of a small collection of cells. Normally one does not get specific recollections by stimulating small areas of brain. The reason for this may be that matters are so delicate that one cannot expect to meet through crude stimulation the conditions for releasing a delicate chain of associations. Consider the situation vis-à-vis the discoveries of Hubel (this volume) on the cat brain. Diffuse stimulation of the retina yields response from only a few of the higher cells. The specific conditions, e.g., for exciting an edge-direction cell, are fairly stringent. In the case that there are cell-assemblies representing notions even more abstract than a disembodied edge-direction, we might expect that the conditions for their excitation would be even more exacting. They might require irregularly-spaced sequences simultaneously arriving at different parts of the cell-assembly. Indeed, one can be sure that the firing conditions for memory associations are highly specific, else we would be in the unhappy condition of associating everything with everything else, or recollect-

ing too many things at any time.

If stimulation by a single electrode is not likely to work, what can we do? As an amateur, I can propose here what seems obvious without the inhibition of less obvious difficulties. We could implant small multiple electrodes and ask a subject to think of various things. We record the activity patterns, repeat the suggestions, and use an on-line computer to try to discover sub-patterns that are correlated with the different topics. Finally we reverse the situation and attempt to stimulate with the discovered patterns, using the same electrodes, while the subject announces his associations. His output is later analysed to see if there is any correlation between his associations and the stimulus classes. Given even the slightest correlation, we would have a hint about how to refine our pattern-analysis procedure to discover something about the requisite stimulation patterns. How does one analyse patterns when one does not know what is being sought? For one thing, we will always have some conjectures with which to bias the analysis program. For another, we might even be able to use some of the new pattern-analysis techniques that are becoming available with the growth of heuristic programming and perhaps even the work in artificial neural nets.

c. The concreteness of abstraction

Our proposal, following Hebb, is that it seems reasonable to expect that rather abstract mental events are represented by the activity of fairly definite groupings of neural activity. An extreme possibility is that one might discover, for each word in one's active vocabulary, a definite group of cells or cell-assembly, and that some such sites could be discovered by computer analysis of multiple electrode activity. At each stage of abstraction there could be another assembly, excited by certain patterns within the assemblies associated with the symbols that combine to form the new stage. We need not suppose anything like a neat hierarchy of abstraction; indeed, this is probably incompatible with a flexible association system. But it does seem necessary to suppose that at each stage the notions must become crystallized, through local decisions. Else there seems to remain little ground for the manipulation of symbolic quantities that seem introspectively to be required. (See again Newell's discussion in this volume.) The separation or "localization" of the symbol-representing activity need not conform to a spatially compact structure, for the "cell assemblies" could include long fibres, or could represent resonant modes of somewhat large structures. The isolation presumably required between assemblies need not be provided by geometric boundaries, but might rather depend on cross-inhibitory mechanisms. Perhaps even the picture of a cell-assembly as a group of functionally-connected neurons is wrong; there are other possibilities. For example, an abstract event might be represented by the route followed by a certain temporal sequence through a tree-like structure, as might be suggested by the learning network model of Feigenbaum and Simon²⁷. There are many possibilities. The important thing is that these must be explored before we rule out the idea that mental events, memories in particular, have rather concrete representations.

REFERENCES

1. RASHEVSKY, N.: *Mathematical Biophysics*, University of Chicago Press, Chicago, Ill.: 1938. (Available as a paperback, Dover Publications, Inc., New York, N. Y., 2 vols., 1960).
2. HOUSEHOLDER, A., and LANDAHL, H. D.: *Mathematical Biophysics of the Central Nervous System*, Principia Press, Bloomington, Ind.: 1945.
3. MCCULLOCH, W. S., and PITTS, W.: A logical calculus of the ideas immanent in nervous activity, *Bull. Math. Biophys.*, 5, 115-137, 1943.
4. KLEENE, S. C.: Representation of events in nerve nets and finite automata, in *Automata Studies*, C. E. Shannon and J. McCarthy, Eds., in *Annals of Mathematics Studies*, no. 34, Princeton University Press, Princeton, N. J., pp. 3-41, 1956.

5. COPI, I. M., ELGOT, C. C., and WRIGHT, J. B.: Realization of events by logical nets, *J. ACM*, 5, 181-196 April, 1958.
6. NEUMANN, J. VON: Probabilistic logics and the synthesis of reliable organisms from unreliable components, in *Automata Studies*, C. E. Shannon and J. McCarthy, Eds., in *Annals of Mathematics Studies*, no. 34, Princeton University Press, Princeton, N. J., pp. 43-98; 1956.
7. McCULLOCH, W. S.: Agathe Tyche,—of nervous nets—the lucky reckoners, *Proc. Symp. on Mechanisation of Thought Processes*, D. V. Blake and A. M. Uttley, Eds., Natl. Phys. Lab., Teddington, Eng., Her Majesty's Stationery Office, London, 2 vols., pp. 613-625; 1959.
8. HEBB, D. O.: *The Organization of Behavior*, John Wiley and Sons, Inc., New York, N. Y., 1949.
9. MINSKY, M. L.: *Neural Nets and the Brain Model Problem*, Ph. D. dissertation, Princeton Univ., Princeton, N. J.; 1954. (Available from University Microfilms, Ann Arbor, Mich.)
10. FARLEY, B. G., and CLARK, W. A.: Simulation of self-organizing system by a digital computer, *IRE Trans. on Information Theory*, vol. PGIT-4, pp. 76-84; September, 1954.
11. CLARK, W. A., and FARLEY, B. G.: Generalization of pattern-recognition in a self-organizing system, *Proc. WJCC*, Session on Learning Machines, W. H. Ware, Chairman, p. 86; 1955.
- 12a. BEURLE, R. L.: Properties of a mass of cells capable of regenerating pulses, *Phil. Trans. of the Royal Society of London*, Series B, 240, 55, 1956.
- 12b. BEURLE, R. L.: Storage and manipulation of information in random networks, in *Aspects of the Theory of Artificial Intelligence*, C. A. Muses, Ed., Plenum Press, N.Y., pp. 19-42; 1962.
- 13a. FARLEY, B. G., and CLARK, W. A.: Activity in networks of neuronlike elements, *Proc. Fourth London Symp. on Information Theory*, C. Cherry, Ed., Butterworth, Ltd., London, Eng.; 1961. Also Academic Press, Inc., New York, N. Y.; 1961.
- 13b. FARLEY, B. G., and CLARK, W. A.: in *Proc. of the 1963 BIONICS Conference*, Wright Air Development Center, to be published.
14. SELFRIDGE, O. G.: Some notes on the theory of flutter, *Arch. Inst. Cardiol. Mex.*, 18, no. 2, 177-187, April, 1948.
15. ROCHESTER, N., HOLLAND, J. H., HAIBT, L. H., and DUDA, W. L.: Test on a cell assembly theory of the action of the brain, using a large digital computer, *IRE Trans. on Information Theory*, IT-2, no.3, 80-93; September 1956.
16. MILNER, P. M.: The cell-assembly; mark II, *Psychol. Rev.*, 64, 242, 1957.
- 16a. MILNER, P. M.: Learning in neural systems, in *Self-Organizing Systems*, M. Yovitts, S. Cameron, Eds., Pergamon Press, New York, N. Y.; 1960.
17. ROSENBLATT, F.: The perceptron: a probabilistic model for information storage and organization in the brain, *Psychol. Rev.*, 65, 386-407; November 1958.
18. ROSENBLATT, F.: *Principles of Neurodynamics*, Spartan Books; 1962.
19. SELFRIDGE, O. G.: Pattern recognition and learning, in *Information Theory, Proc. Third London Symp. on Information Theory*, C. Cherry, Ed., Thornton Butterworth, Ltd., London, Eng., p. 345, 1956.
- 19a. SELFRIDGE, O. G.: Pandemonium: a paradigm for learning, in *Proc. Symp. on Mechanisation of Thought Processes*, Natl. Phys. Lab., Teddington, Eng., Her Majesty's Stationery Office, London, 2 vols., pp. 511-529; 1959.
20. MINSKY, M. L., and SELFRIDGE, O. G.: Learning in random nets, *Proc. Fourth London Symp. on Information Theory*, C. Cherry, Ed., Butterworth, Ltd., London, Eng.; 1961.
21. UTTLEY, A. M.: Conditional probability computing in a nervous system, in *Proc. Symp. on Mechanisation of Thought Processes*, D. V. Blake and A. M. Uttley, Eds., Natl. Phys. Lab., Teddington, Eng., Her Majesty's Stationery Office, London, 2 vols., pp. 121-146; 1959.
22. GAMBA, A.: *Proc. IRE*, 49, 349, 1961.
- 22a. GAMBA, A.: *Suppl. Nuovo Cimento*, 23, 146, 1961.
23. BRYAN, J.: *Proc. of the 1963 BIONICS Conference*, Wright Air Development Center, to be published.
24. MINSKY, M. L.: Steps toward artificial intelligence, in Special Computer Issue, *Proc. IRE*.

**CS-TR Scanning Project
Document Control Form**

Date : 11/30/95

Report # AIM - 48

Each of the following should be identified by a checkmark:
Originating Department:

- ☒ Artificial Intelligence Laboratory (AI)
☐ Laboratory for Computer Science (LCS)

Document Type:

- ☐ Technical Report (TR) ☒ Technical Memo (TM)
☐ Other: _____

Document Information

Number of pages: 10 (14 IMAGES)
Not to include DOD forms, printer instructions, etc... original pages only.

Originals are:

☐ Single-sided or

☒ Double-sided

Intended to be printed as :

☐ Single-sided or

☒ Double-sided

Print type:

- ☐ Typewriter ☐ Offset Press ☐ Laser Print
☐ InkJet Printer ☐ Unknown ☒ Other: COPY OF JOURNAL ARTICLE

Check each if included with document:

- ☐ DOD Form ☐ Funding Agent Form ☐ Cover Page
☐ Spine ☐ Printers Notes ☐ Photo negatives
☐ Other: _____

Page Data:

Blank Pages (by page number): _____

Photographs/Tonal Material (by page number): _____

Other (note description/page number):

Description :

Page Number:

IMAGE MAP: (1-10) UNHED TITLE & BLANK, 303-310
(11-14) SCAN CONTROL TARGETS (3)

Scanning Agent Signoff:

Date Received: 11/30/95 Date Scanned: 12/11/95

Date Returned: 12/14/95

Scanning Agent Signature: Michael W. Cook

Scanning Agent Identification Target

Scanning of this document was supported in part by the **Corporation for National Research Initiatives**, using funds from the **Advanced Research Projects Agency** of the **United states Government** under Grant: **MDA972-92-J1029**.

The scanning agent for this project was the **Document Services** department of the **M.I.T Libraries**. Technical support for this project was also provided by the **M.I.T. Laboratory for Computer Sciences**.

