AD-A154 507 THE REPRESENTATION OF KNOWLEDGE IN IMAGE UNDERSTANDING 1/1
(U) MASSACHUSETTS UNIV AMHERST DEPT OF COMPUTER AND
INFORMATION SCIENCE D N SPINELLI MAR 85
UNCLASSIFIED AFOSR-TR-85-0448 F/G 6/4 NL

END
FILMED
DTIC

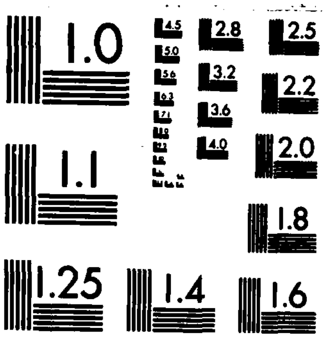
1.0
4.5
5.0
5.6
6.3
7.1
2.8
3.2
3.6
4.0
2.5
2.2
2.0
1.1
1.8
1.25
1.4
1.6

②

Annual Scientific Report
March 1985
AFOSR-83-0207

AD-A154 507

THE REPRESENTATION OF KNOWLEDGE IN IMAGE UNDERSTANDING

on For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A/1	

Computer and Information Science
University of Massachusetts
Amherst, MA 01003

Dr. D. N. Spinelli

DTIC
COPY
INSPECTED
3

DTIC
ELECTE
JUN 4 1985
S D
E

Controlling Office: Air Force Office of Scientific Research/NL
Bolling AFB, DC 20332-6448

DTIC FILE COPY

85 5 07 098

REPRODUCED AT GOVERNMENT EXPENSE

UNCLASSIFIED

ECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
DECLASSIFICATION/DOWNGRADING SCHEDULE	
PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR- 85-0448
a. NAME OF PERFORMING ORGANIZATION Computer and Information Science 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research/NL
c. ADDRESS (City, State and ZIP Code) University of Massachusetts at Amherst Amherst, MA 01003	7b. ADDRESS (City, State and ZIP Code) Building 410 Bolling AFB, DC 20332-6448
a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR 8b. OFFICE SYMBOL (If applicable) NL	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-83-0207
ADDRESS (City, State and ZIP Code) Building 410 Bolling AFB DC 20332-6448	10. SOURCE OF FUNDING NOS. PROGRAM ELEMENT NO. 61102F PROJECT NO. 2312 TASK NO. A1 WORK UNIT NO.

1. TITLE (Include Security Classification)
The Representation of Knowledge in Image Understanding

2. PERSONAL AUTHOR(S)
Dr. D. N. Spinelli

3a. TYPE OF REPORT Annual Scientific	13b. TIME COVERED FROM 5/1/83 TO 4/30/84	14. DATE OF REPORT (Yr., Mo., Day) March 1985	15. PAGE COUNT 17

3. SUPPLEMENTARY NOTATION

COSATI CODES		
FIELD	GROUP	SUB. GR.

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
neuroscience, image understanding, adaptive architecture

). ABSTRACT (Continue on reverse if necessary and identify by block number)

Image understanding in animal brains is subserved by mechanisms that exhibit massively parallel and adaptive architecture. In the higher vertebrates, such as cats, that have frontally located eyes, relatively simple and well understood retinal and geniculate structures transform the image into a firing pattern of optic radiation fibers that feed into several areas of cortex dedicated to vision. We have shown that neurons in these areas are very adaptive. In fact, given the proper conditions, it can be demonstrated that the shape of their receptive fields comes to resemble the shape of the visual patterns viewed by the animal. To our knowledge this is the only instance in which recording from brains has produced recognizable images of what an animal has seen in the past. The practical and theoretical implications are staggering. We have created a visual experience which has unique properties, so that when we find its traces in the brain we have no

). DISTRIBUTION/AVAILABILITY OF ABSTRACT NCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS ☐	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE NUMBER (Include Area Code) (202) 767-5021 22c. OFFICE SYMBOL NL

D FORM 1473, 83 APR   EDITION OF 1 JAN 73 IS OBSOLETE.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

question that we are dealing with adaptation and memory. Quite simply, the experience consists of having the animal view patterns which are different for the two eyes. This does not occur in nature for animals with frontally located eyes: neurons that show receptive fields which reflect this experience then unequivocally have adapted and belong to those neural nets that record experience. We believe that understanding adaptation and knowledge representation is fundamental to make progress in image understanding by animal brains. We have been studying the anatomical structures and the temporal requirements that lead to adaptation. In this report we make a brief excursion into some experiments which have been demonstrated in the past. We also discuss the enormous complexity of the problem at hand and we conclude that only a method which allows the reading out of recognizable visual memories has any chance to make progress in this complex endeavor. We have determined some of the parameters that produce powerful adaptation and have methods that allow memory read-outs. A preliminary conclusion is that image understanding requires a learning principle that takes into account the nature of the information and not just temporal and/or spatial relationships. Animal brains possess extremely effective vision systems. Architectural and functional principles gained by studying them will certainly lead to new ideas for new computer architectures especially in the fields of machine vision, adaptation, and parallel computation.

(2)

THE REPRESENTATION OF KNOWLEDGE IN IMAGE UNDERSTANDING

Annual Technical Report: AFOSR-83-0207

Period: 5/1/83-4/30/84

Principal Investigator: D. N. Spinelli

ABSTRACT

Image understanding in animal brains is subserved by mechanisms that exhibit massively parallel and adaptive architecture. In the higher vertebrates, such as cats, that have frontally located eyes, relatively simple and well understood retinal and geniculate structures transform the image into a firing pattern of optic radiation fibers that feed into several areas of cortex dedicated to vision. We have shown that neurons in these areas are very adaptive. In fact, given the proper conditions, it can be demonstrated that the shape of their receptive fields comes to resemble the shape of the visual patterns viewed by the animal. To our knowledge this is the only instance in which recording from brains has produced recognizable images of what an animal has seen in the past. The practical and theoretical implications are staggering. We have created a visual experience which has unique properties, so that when we find its traces in the brain we have no question that we are dealing with adaptation and memory. Quite simply, the experience consists of having the animal view patterns which are different for the two eyes. This does not occur in nature for animals with frontally located eyes: neurons that show receptive fields which reflect this experience then unequivocally have adapted and belong to those neural nets that record experience. The author ~~We~~ believes that understanding adaptation and knowledge representation is fundamental to make progress in image understanding by animal brains. He ~~We have~~ has been studying the anatomical structures and the temporal requirements that lead to adaptation. In this report we makes a brief excursion into some experiments which have been demonstrated in the past. ~~We~~ also discussed is the enormous complexity of the problem at hand and it is ~~we~~ concluded that only a method which allows the reading out of recognizable visual memories has any chance to make progress in this complex endeavor. ~~We have~~ determined were some of the parameters that produce powerful adaptation and have methods that allow memory read-outs. A preliminary conclusion is that image understanding requires a learning principle that takes into account the nature of the information and not just temporal and/or spatial relationships. Animal brains possess extremely effective vision systems. Architectural and functional principles gained by studying them will certainly lead to new ideas for new computer architectures especially in the fields of machine vision, adaptation, and parallel computation.

THE REPRESENTATION OF KNOWLEDGE IN IMAGE UNDERSTANDING

Annual Technical Report: AFOSR-83-0207

Period: 5/1/83-4/30/84

Principal Investigator: D. N. Spinelli

Image understanding is an extremely complex task that is performed with astonishing ease and speed by animal brains. The human retina contains about 10^7 sensors and the optic nerve about 1.5×10^6 fibers. Given that the maximum frequency of firing is about 1,500/sec it is quite probable that the equivalent of 2.25 gigabits per second are sent to the brain by one eye alone. Given the very slow response time of neurons one must assume that massive parallelism and extremely effective ways of representing knowledge are at the core of the brain's image understanding processes. Much is now known as to the kind of processing that is performed at the retina, lateral geniculate body and visual cortex, but a great deal remains to be done if we want to understand how visual information is stored. What this means is locating in the brain neurophysiological changes which are <u>unequivocally</u> produced by experience (visual or otherwise). This is an extremely important goal in neuroscience because locating such traces, understanding the mechanisms that bring them about, and the nature of their organization would simultaneously gain us tremendous insights into the nature of learning, the encoding of sensory information, the representation of knowledge and how that knowledge is accessed and used. Further, we might gain some ideas as to the kind of architecture that makes brains so fast in spite of the slowness of their components and so resistent to damage.

The search for the memory trace has been on for a few decades now, at various levels in the hierarchy of the brain. This hierarchy is quite unlike that of a computer, however a simple analogy with computers can help in identifying these levels. Memory, in a computer, can be studied at the level of a single bit, or at the level of bytes and words or from the point of view of one of the many types of data strucures, such as arrays, trees, graphs, etc. that computer scientists use.

Neurophysiologically we can look for traces of experience at the level of a single synapse, a single neuron, or networks of neurons. These traces must bear a recognizable, and unequivocally so, relationship to the experience that induced them.

We have developed a methodology, and obtained results, that we firmly believe has enabled us to locate traces of visual memories and identify some of the elements that make an experience memorable, unique, and effective so that it will have a large and recognizable effect on neural structures and function. In fact, we can read out a visual experience from a single neuron! We further believe that the study of visual memories will greatly increase our understanding of how brains process images and build effective knowledge representations that enable the humblest of animals to navigate, forage, and survive in complex environments.

What has neuroscience achieved in the search for memory traces?

It's been quite a few years since Lashley wrote his famous paper "Search for the Engram." That paper pinpointed the fundamental place that memory, and knowledge representation has in the functioning of the brain if we are to understand behavior and mentation. A generation of "engram hunters," mostly empirical neurobiologists guided by their own often powerful working hypotheses, ensued and has continued that search with impressive results.

There is, however, substantive disagreement as to "what" constitutes a memory trace. Even the "where," that is locating memory traces, seems to have no foolproof methodology.

The persistence of these problems, at a time when the empirical field is at its highest level of hit rate in terms of locating and describing traces of "experience" in the brain, seems due to the lack of a clear set of principles by which investigators can recognize and classify the different tiers in the hierarchical structural and functional organization of what we call memory.

Theoreticians have also not provided a list of possible memory structures accompanied by manageable verification tests.

What is memory? Where is it? These are problems of inordinate complexity. The principal problem seems to be that of recognition, that is, given that one encounters a memory in the brain, how does one recognize it for what it is? Sokolov's idea (heard at the Asilomar conference) is that unless some kind of isomorphism exists between what is encountered and what is to be remembered we would never know. Biologists in general and neurobiologists in particular have always turned to technology as a source for inspiration and modelling. Even such a "simple" system such as blood pumping which had been tackled and almost understood by Leonardo's powerful mind, had to wait for the mechanical pump to be completely unraveled. Understanding gained in one field is used analogically to crack problems in other fields. Not too long ago it was the telephone switchboard that inspired memory theories, now it's the computer. In both cases one is dealing with structures that handle, route and store information. It is quite possible that other machines will provide inspiration at some future time or that the direction of flow might change, as it often does, but at the moment it's in the field of computer science that one finds the most detailed and powerful realizations that can help, analogically, to build a base from which it will be possible to further our understanding of biological memories. In any case the study of different systems might begin to point the way to higher level abstractions useful or even necessary to chart the way to further progress.

What is an Engram? (Some Informal Considerations)

Memory, viewed as traces of experience, is necessarily massively distributed. As an example of this, let's consider a tennis player as an organism trained to recognize a certain spherical object capable of high velocity movement, to intercept it and strike it with another object in such a way that its path will take a desired course. Many structures have been permanently or semi-permanently changed. Macroscopically we see that the muscles of the forearm used in the game are hypertrophic, we know that has been brought about by lots of activity which means that large amounts of acethil choline have been released at the neuromuscular junctions over a substantive period of time. As we know that the metabolic machinery for the end plate is in the soma of the motoneuron we infer that, if we did a histological preparation of the spinal cord of the player, we would see a clear hypertrophy of the cells innervating the tennis arm as compared to the other one. It wouldn't be surprising at all if this asymmetry were to reach into higher and higher levels of the neuraxis, after all we do know that activity leads to trophism whereas inactivity or non-use to atrophy. One begins to doubt at this point whether increase in strength of a synapse during or after a training procedure is sufficient criterion to show memory traces. In fact, let's do a thought experiment. Imagine that we have a reduced system which contains an alpha neuron pool, a gamma pool, and input-output nerves from a muscle. If we pull on the muscle the negative feedback will be activated and the muscle will contract. Let's do this rhythmically over a sufficient number of days and the muscle will grow stronger, thereby opposing the stretch more effectively. We could surmise that the alpha-gamma loop has learned to oppose the movement more effectively and it has, but few would call it true learning.

Going back to our tennis player, we would probably find differences in the dendritic arborizations of cells in the cerebellum, because tennis involves a lot of skilled movements. Sensory motor cortex would be definitely involved and we would expect the cortical representations for the tennis and non-tennis arm to be different macroscopically, microscopically, and at the level of cellular responsivity.

Now, just for the moment, posit a simpler system, that is, a robot built to play tennis. First of all, it's clear that the arm selected to play tennis needs to have faster actuators and more powerful motors than the one that doesn't. (We could build it perfectly symmetric at extra cost naturally.) Then the transistors that drive the current in the motors will have to be physically larger to handle the higher wattage, wires thicker, etc. In RAM and/or ROM larger areas will be occupied by routines specifically constructed to handle arm movements that have to do with the game than for the other arm which will have none of that. There will be large areas dealing with running, balance, etc. which literally have nothing to do with the game itself, but without which the game could not be played. Finally, specific image processing procedures occupying RAM space roughly related to their complexity will also exist. Because it's easier to do it that way chances are that "clusters" of like routines and data will exist with a few things scattered erratically here and there because they were added or patched in later. If we want to be precise we could identify areas containing "procedural" knowledge (procedures) and areas containing "declarative" knowledge (arrays, trees, lists . . . that is one or more of those data structures that programmers are familiar with).

In designing such a system or any other complex program for that matter, a programmer would use a top-down/bottom-up approach. That is, the task would be broken into subtasks, the subtasks into sub-sub tasks, etc. until small enough processes were identified. At this point the bottom-up part would begin, our programmer would look up the system library to see which elementary processes are already available in the system (just as a human tennis player does not have to learn to walk up to a ball but makes use of a walking skill avalable from previous learning) and code only procedures that do not exist already. We see here that a larger system will have a larger library, making less original code necessary. Also that when a system routine is called and parameters passed, the call itself will give no indication of the complexity of the process called. As the new program comes into being it will contain procedures which might be in turn useful to and will be called by subsequently acquired schemas, e.g., a badminton one. Thus even with a memory map we would rediscover mass action by making a lesion in the domain-specific tennis area! That is because all schemas learned after the tennis one that call tennis routines would be impaired.

Where is the memory for tennis in our robot? Imagine doing a study of the robot's "brain." First we would have to locate memory. But that would be only a small beginning. We really have no way to identify those procedures that recognize the ball, or those that detect contact with the racket, etc. We could design cogent experiments, because we know the program, the data structures, etc., thus we could formulate questions to locate and identify what we know has to be there somewhere.

"Selective lesions" done by changing all instructions in a procedure to no-ops will impair performance almost anyplace. If our programmer was clever most procedures will be general, that is, invocable by other programs that have nothing to do with tennis; conversely, the tennis program might invoke system routines.

I guess that one could more properly say that our robot has been given (has learned) a tennis schema and that would comprise all of those changes and additions that were necessary in the hardware and software to enable it to play the game. These will consist of domain-specific new procedures, calls to (with appropriate parameters) already present procedures, plus additions and changes to the hardware. Even remote systems, such as temperature control, might have been strengthened. Undoubtedly a "sedentary" robot will not warrant the same expense as an athletic one! Where is the engram in the schema? One possibility is to call the whole schema for tennis the tennis engram. Somehow even though the whole schema is necessary this seems unsatisfactory. The schema engages too much of the system. Domain specific knowledge seems a better candidate. As neurobiologists we would like to find a place that, when lesioned, impairs or abolishes tennis playing only. However, changing only a few instructions (we know) can completely incapacitate large programs. Leaving method undefined for the moment, it would seem that removal of all the code that has an effect on tennis and nothing else would fit the requirement of identifying the locus of the tennis engram.

Alternatively, if we could "read out" programs and data structures we would have little doubt as to what is in fact the "tennis program." The above rather lengthy considerations convince us that even a computer scientist could not identify memories, that is, programs and data in a computer unless said memories could be read out in a "readable," that is, task-related way. These problems are even more complex in natural brains.

Because of this massive distribution of permanent changes many neurobiologists prefer to use the expression "traces of experience" rather than engram or memory trace even though these labels are reappearing more frequently lately.

One of the tasks of this paper will be to attempt to identify some criteria by which some traces of experience can legitimately be called memory traces or at least unambiguously determine to which level in the hierarchy of traces of experience a given trace belongs. This is not just a naming problem. Our claim is that in-depth study of memory traces will immeasurably further our understanding of brain; on the other hand most "traces of experience" could prove to be relatively less important. For example, self-adjusting feed-back loops (much as some cars have self-adjusting brakes) or feed-forward ones to take care of local changes such as arm weight, image slippage on visual tracking due to imperfect compensation of head and body movement, etc.

Methods Used in the Search for the Engram

In general neurobiologists have used the following classes of approaches:

1) Aimed at LOCATING a trace by lesioning, cooling or chemically treating brain areas of various size. It's a bit unclear if the community would consider such a localization "finding the engram" or simply a prodromal step which is necessary to the actual detailed study of neural circuits and their logic (what Hubel and Wiesel have referred to as the functional architecture). While this achievement would be substantive (a bit like proving that genetic information exists and is located in the chromosomes) it would tell us nothing as to how it all works . . . from the point of view of

computer science it is analogous to discovering where RAM is, i.e., no information as to circuitry, how to read or store, nature of the encoding/decoding process, distinction between areas that contain procedural vs. declarative knowledge, etc. This approach has the appeal of technical simplicity, but is also easily criticizable. A number of investigators have produced fascinating results in this area, i.e., Sperry, Thompson, etc.

2) Another approach has been the identification of the electrical correlates of memory. The method is based on the very reasonable assumption that whatever the ultimate trace, neural activity must be present every time a "memory" arrives or departs (memory read-in, read-out). Macro, semi-micro and microelectrodes have been used to detect these potentials aided by a variety of statistical techniques aimed at demonstrating correlations between electrical activity, brain regions, and/or behavior and memory. A large body of impressive results exist in this category. One of the earliest representatives among investigators is R. John Killam followed by others such as Spinelli, Pribram, and Kamback. Problems with the method are that at best it provides localization, but no circuitry. Good things are that in the hands of investigators that pay close attention to time it can provide a rough block chart of the system.

J. Olds, who came to an untimely death in the midst of a brilliant career, was looking for memory correlates in the hippocampus using chronic microelectrodes, his criterion and goal was to locate the earliest locus of correlation between all activity and behavior.

3) The chemistry of memory is also very important, but we will not be discussing it as it is comparable to the study of solid state physics. Computers as we know them could not exist without it, but it contributes a different kind of understanding than the one we are seeking.

4) The FUNCTIONAL ANATOMY approach, spearheaded by Hubel and Wiesel who have not really been in the business of locating memory traces, but who have followed an impeccable methodology based on the study of single cells and their receptive fields and the anatomical foundations of columnar organization. This method, possibly limited to vision, is in our opinion the only one that at present seems capable of satisfying Sokolov isomorphic recognition criterion. Combining it with careful measures of delay, as Stone and Hoffman have done, yields a comprehensive approach to the study of neural circuitry and possible plastic phenomena.

Given these considerations, let's review a number of empirical findings that appear to qualify as candidates for a place in one or more of the tiers of the functional hierarchy of memory.

It seems only fair to begin with the fundamental work of Lloyd on postetanic potentiation. In a series of classic studies Lloyd demonstrated in 1947-58 that the responsivity of synapses depended on preceding near term (up to 10 mins) past activity that had impinged on that synapse (Lloyd's experiments were actually done on motorneuron pools and took advantage of the subliminal fringe). Specifically, he demonstrated that following stimulation at 2-300/sec the monosynaptic reflex recorded from a ventral root (as a response to .1/sec stimulation of the corresponding dorsal root) was first deeply depressed, then greatly magnified for many minutes. This temporary modification could also be produced by natural activation of the input pathways and Lloyd convincingly proved that hyperpolarization and/or depolarization of the presynaptic terminals (leading to increased or decreased transmitter release of transmitter) was involved.

Later on McIntyre (1953) demonstrated that the decay constant of postetanic potentiation could be lengthened to over 30 mins if all activity on the path had been prevented for a period of 30 days. These experiments and a host of others convincingly show that permanent or semipermanent synaptic modifications can be demonstrated at the spinal cord level and raise the following question: is modifiability of neural connectivity a property of specialized junctions? Or is it a general property of neural tissue in general? The two underlying hypotheses are not mutually exclusive and chances are that both are at work.

To my knowledge none of these investigators attempted to identify possible anatomical concomitants of these effects. Studies of this type correspond, in our metaphor, to turning a bit on, establishing a connection, or a flag, or from a structural standpoint to a functional description of control gates in transistors, some of which have a fleeting action whereas some others, i.e., in EPROMS, have decay times measurable in years. The technology of some EPROMS require that a special control wire be active for storage to take place. It is interesting to note that neurobiologists are still divided on this issue—some claiming that temporal correlation between events suffices to create a link (Hebb), while others maintain that stimuli with special status, that is unconditioned reinforcers (Ito, Barto and Sutton), are necessary for storage to take place. Irrespective of the ultimate mechanics, however, we are dealing with bit setting or single gate links (possibly analog in nature like a valve), and from our previous discussion we would expect large areas of the brain to exhibit bit setting when any task of even moderate complexity is learned. And that is why Lashley generated the theory of mass action, just as he would have if he had carried out his investigation in our tennis playing robot. Experiments which explicitly or implicitly aim at bit setting are appropriate to the study of synaptic events (Kandel) and fundamentally so. However, this methodoogy is, in our opinion, unsuited for studies that aim at localization of the memory trace and the type of data structures involved (array, holographic, graph, etc.) simply because the most trivial of tasks is likely to produce bit settings all over the place (e.g., if the required response is to extend the left arm forward there will be equilibrating contractions in much of the rest of the body).

Of all the methods used by neurobiologists, localization by lesion has been the most widely used. It is also the only method that, partly because of gunshot wounds to the head or other forms of head trauma, has generated extensive data from human brains.

It's definitely too large a body of data to be reviewed comprehensively. In essence, there are two facets to it:

1) Extensive ablation of brain structures that do not contain the engram so as to leave untouched that smallest brain piece capable of supporting the behavior—Sperry's cortical island technique.

2) Minimal ablation of a brain structure capable of abolishing the behavior—most widely used and fraught with underlying assumptions. Here impressive results have been achieved by K. H. Pribram on IT cortex using minimum overlap from several lesions, by C. Butter on IT also using the same method, and more recently by R. Thompson with small lesions in the cerebellum. The first set of experiments addresses itself to the image understanding part of the schema, whereas the second one seems to impact on the "skilled" movement part of it.

We are all familiar, of course, with the search for the engram. What is often forgotten, or suppressed in our minds, is that the search has in fact been quite successful! In my opinion, a trace of memory was first located by R. Sperry (1961) in his cortical island experiment. A monkey was taught a behavior involving a hand, and then the cortex of the contralateral hemisphere was extensively ablated while sparing the hand representation: The behavior remained. Later lesioning of the cortical island abolished the behavior. Sperry argued persuasively that memory for the behavior had to be in the cortical island. Surprisingly, the neuroscience community did not react with a Manhattan Project.

Another trace of memory, or at least its electrical correlates, was captured by E. R. John (John and Killam, 1959) using his creative idea of tracer stimuli and later making imaginative use of averaged evoked potentials and sophisticated mathematical analysis. Massive neural traces of memory were found by Wiesel and Hubel (1965) in their monocular closure experiment; here for the first time, one could reliably produce massive neural changes, demonstrable at the single cell level, using only experience.

Van Harreveld and Fifkova (n.d.) demonstrated the relation of spine swelling in the hippocampus to repeated stimulation of afferents, thus providing one of the first demonstrations that neural activity results in structural changes. I cannot resist the temptation to mention my own work (Hirsch and Spinelli, 1970, 1971; Spinelli, Hirsch, Phelps, and Metzler, 1972) in which we actually managed to produce, while recording from single cells, a recognizable picture of a simple pattern (three bars) that a cat had seen up to a year before. Later, a trace of memory was again, in my opinion, conclusively demonstrated in a simple experiment in which a simple behavior, unique in its structure, engendered cells in visual and motor cortex tuned to features in a way never encountered in "normal" cats (Spinelli and Jensen, 1979).

The above experiments have shown what Lashley (1950) could not show. We have gained over the years the ability to study the properties of a single cell quite precisely; this double-edged skill has enabled us to provide, yet limited us to, remarkable but atomic demonstrations of traces of memory. Scientists versed in computer science can see right away that future progress depends critically on a comprehensive theory that is rich and precise enough to be tested, that goes beyond bits and pieces, and that addresses itself to the overall organization of memory.

It's because of the above considerations that we have developed a methodology that:

1) Enables us to read out visual memories in a recognizable form by mapping visual receptive fields with a computer.

2) Enables us to be certain that what we see is produced by the experience, because we designed an experience that, while simple, doesn't occur naturally.

3) Enables us to trace neural nets by using horseradish peroxidase. Finally, with the help of heterostatic theory (Klopf, 1982) we have tuned the visual experience so that it is not only most effective, but it's easily related to modern learning theory.

Our experiments have shown that time is a critical element in determining adaptation and that possibly the time function has two peaks, one at .4 secs and one at 4-6 secs. Further, it has become evident that one of the concerns of learning theory, that is, the nature of the information itself, needs to be addressed. It could very well be that what learning theorists call process learning, that is, the learning that establishes the context in which all other learning takes place, is an essential ingredient in image understanding and knowledge representation. These problems await further research.

REFERENCES

Barto, A. G., & Sutton, R. S. (1981) Landmark learning: An illustration of associative search. Biological Cybernetics 42: 1-8.

Dickinson, A. Contemporary Animal Learning Theory.

Klopf, A. H. (1982) The Hedonistic Neuron: A Theory of Memory, Learning, and Intelligence. Washington, DC: Hemisphere. potentiation of monosynaptic reflex pathways of the spinal cord. Journal of General Physiology 33: 147-170.

McIntyre, A. K., & Eccles, J. C. (1953) The effect of disuse and of activity on mammalian spinal reflexes. Journal of Physiology 121: 492-516.

Sperry, R. W. (1961) Cerebral organization and behavior. Science 133: 1749.

Sperry, R. W., & Miller, N. (1955) Pattern perception following insertion of mica plates into visual cortex. Journal of Comparative and Physiological Psychology 48: 463-69.

Spinelli, D. N. (1975) Silver tipped metal microelectrodes: A new method for recording and staining single neurones. Brain Research 91: 271-75.

Spinelli, D. N., Hirsch, H. V. B., Phelps, R. W., & Metzler, J. (1972) Visual experience as a determinant of the response characteristics of cortical receptive fields in cats. Experimental Brain Research 15: 289-304.

Spinelli, D. N., & Jensen, F. E. (1979) Plasticity: The mirror of experience. Science 2023: 75-79.

Van Harreveld, A., & Fifkova, E. (n.d.) Mechanism for potentiation and short term memory. Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen, Series C 78: 21-24.

Wiesel, T. N., & Hubel, D. (1965) Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. Journal of Neurophysiology 28: 1029-40.