ATTENTION AND COMPUTATION IN THE BRAIN

Edward F. Storm

"Each of us at any moment of the waking day is a whole bundle of acts simultaneously proceeding. In no case does any other of all the doings of the moment disturb the one focal doing. . . . Should it do so then the pattern changes and the disturbing piece becomes usually the keypiece of a new pattern which supplants the previous. . . . The individual cannot be the seat of two focal acts at once."

Sir Charles Scott Sherrington
Man On His Nature

Introduction

Neither, I believe, can the individual be the seat of two simultaneous acts of focal awareness. There may be other "doings of the moment" which support the object of focal awareness, but that object is always specific, dominant and unitary in form. On the other hand, the object of awareness changes easily in response to events occurring either within or without the brain. We then ask what kind of organization supports the manipulation of awareness, and we seek specific instruments for this manipulation. I will propose here that the attending mechanism has essential computational characteristics. I will also show that at the most familiar level of structure and function, the brain is not a computer. For this observation to be persuasive, it will be necessary to be quite precise about what a computer is. Fortunately, this precision in definition is available from existing studies in the theory of computation.

The variety of applications of computing machines and the metaphors which they provide for abstract processes suggest that computing may be something more than technological innovation, that one may expect to find computational entities at almost any level of organization of reality. While many processes studied in the physical sciences are best idealized in the form of differential equations whose solutions are functions on continuous domains, there appear to be more abstractly organized activities whose descriptions are best expressed in computational terms. Certain aspects of the structure of human language and of human cognition seem to be organized into finite and discrete categories, and the associated processes are in some cases well imitated by computer programs. At the same time, one may doubt that a complete description of the human personality in important respects can be obtained in computational terms. The issue is to determine the nature of the relation between computation and the human personality. It is therefore both reasonable and sensible to inquire whether there are any computational entities in reality apart from the technological artifacts of human creation and, if there are, to identify them. One may ask, for example, whether interactions among elementary particles have any computational properties. Is the genetic code really a program for a computing process? Is the brain a computer?

In the first section of this study some general remarks provide a perspective for investigating the nature of physical computation. We see, for example, that computation is a physical phenomenon, that it is in principle uninterpreted, and that its relation to causality is problematic. Then in the next section, physical conditions are described which any entity must meet in order to be part of a computation. Then I review the precise definitions for computing and observe that a certain idealized computer must meet the physical conditions specified if it is to behave as defined.

Then follows a brief review of some facts about the membrane properties of neurons and their signal initiation and propagation functions. These structures and functions are seen not to meet the conditions required of computational entities, suggesting the possibility that there are no computational entities of any kind in any biological systems. If this is indeed the case, and if at the same time we believe that our behavior, physical or mental, has computational structure, then we may have to seek the source of biological computation in structures which are not completely characterized in physical terms. For example, if we believe that there are computational structures and processes in the psychological domain, then the preceding observations raise serious difficulties for any psychophysical identity theory. Then I report on my own experience with awareness, focusing on the processes of attending. On the basis of this introspective evidence, I conclude that the form of attending is purely computational, and that the act of attending arises as a result of the emergence of computational forms out of a non-computational substrate. Finally, the implications of this conclusion are briefly discussed.

The General Concept of Computation

Throughout this study, when I speak of a computer I specifically exclude analog devices from the discussion. Complex assemblies of digital and analog devices are powerful instruments, but they are in general different from purely digital machines or purely analog machines. Subsequent investigation may reveal that such complex assemblies are similar to what we find in living systems, but we must first know what constitutes a digital device occurring in a biological system or as an aspect of a biological system, before we can recognize a hybrid.

If we think that brains and their constituents act as mediators between psychological and physical reality, then we will ask whether a particular level of organization of brain structure and function is appropriate for the psychological realities, however these are specified. We may ask whether the physiological correspondents to cognition, if any, are found at the level of elementary particle interactions, at the level of molecular reactions, at the level at which neurons initiate and distribute signals, at the level of brain waves, or at some other as yet unspecified level. If we determine conclusively that no level supports a computational interpretation, then we may have to recognize the importance of essentially abstract organizing principles at the level of thought, or else discard rational categories from our analysis of thought. Here I will only consider whether brain structure and function at the level of the nerve impulse can be interpreted in computational terms.

Four kinds of entities will figure in this discussion—programs or algorithms, structures, computers and computations. In informal terms, a program is a complete specification of what is required with respect to some particular computational task. A structure is a systematic formal representation of what is to be manipulated by the computer. The computer itself is a physical device capable of "reading" structures, of interpreting programs and of committing the acts specified by them on the structures. A computation is a physical event realizing one complete act committed on some structure. I will use the phrase "computing agent" to refer either to an existing artifact or to a hypothetical biological computer. The distinction is real since it has not yet been shown that there are any biological computers.

All aspects of a computer's structure and function are both finite and discrete. A computer is finite at the abstract level, in that it may not have infinitely many parts, and may not contemplate infinitely many elementary things. At the concrete level it is, like any other physical object we apprehend, finite in spatial and temporal extent. A discrete or digital character follows from the abstract finitude, but is important enough to receive separate mention. And the mathematical properties of denseness and compactness are excluded from the irreducible objects which computers manipulate. Although computations may be described in abstract terms, a computation itself, as opposed to a description of it or a representation of it, is a genuine physical event involving the time-ordered manipulation of matter and energy in orderly ways. We recognize then that there may be universal principles of physical computing, just as we think there are universal principles of gravitation or of electromagnetic radiation. We may discover that only certain aspects of a biological event are computational. And it may be important to distinguish this phenomenon from that in which physical systems "approximate" computers, itself a very unusual notion. In any case, a computation is a physical process.

Another important observation is that computation consists in the manipulation of pure form. Once we understand this we can appreciate that a computer does not "do" arithmetic. It manipulates forms in such a way that we may reliably interpret those manipulations as arithmetic operations. Specific logical circuits may be designed with arithmetic in mind, but it is entirely possible to construct a digital computer which has no explicit arithmetical capability, just as one could be constructed with no explicit capability to manipulate alphabetic characters. This fact about computation has two important consequences. The first is that if computation occurs in physical reality, then that reality must be able to assume pure form in such a way that it can be computationally manipulated. The second is that computations must in principle be describable entirely without reference to interpretations. If computations are to have substance as well as form, then it will be essential to have a precise specification of what that substance is, what hidden variables determine its character, and some kind of guarantee that it is something other than a selected aspect of form rearranged or disguised. In the same way, computational entities may not be distinguished solely on the basis of interpretation. All general purpose computers can be made to do arithmetic either directly or indirectly, but the ability to do arithmetic does not guarantee that an object is a computer.

In this sense computers do not "play" chess, translate or understand natural languages, identify spectral lines or simulate neural nets. A computer engages in the manipulation of form, a manipulation that we may choose to interpret in a way. In the ordinary situation there is a prior specification of certain complex forms, a specification of complex acts to be committed upon these forms, and a convenient interpretation which gives the computations significance but never directly influences them. It is really not remarkable that computers can be made to do all this complex form manipulation. The miracle is that matter organized out of organic molecules and living cells has conceived and become conscious of the power of the manipulations of form in the first place.

An even more important issue is that from a theoretical point of view a computer's actions are not causal, at least in any familiar sense of the word. Rather, a computer administers the form of causally described events. A computational scheme specifies that certain pairs of forms are related in such a way that the computer replaces the first form with the second. The persistent iteration of such changes in form is part of the essence of computation. It is difficult to appreciate what kind of causal principle could account for such transformations, whose subject matter is devoid of interpretation. But the fact is that biological systems seem to achieve this manipulation of form and to appreciate that the manipulation of form is an important aspect of physical reality.

Finally, it should be noted that there is no vagueness or indefiniteness about what a computer is. Differing formulations of one or another aspect of the concepts of computation were presented in the early and mid-nineteen thirties, formulations which have been shown to be intimately related, through a well-defined sense of equivalence, and which have stood the test of time in two important respects. The particular and detailed issues associated with computation were specified in those formulations, and these issues are in principle still the foundation of computational practice. And each of these formulations was seen to treat very broad and general issues in a uniform way. Whatever could be achieved within one formulation could be achieved in an exactly specified way in any other formulation. In short, we have exact specifications for what constitutes a computer, an algorithm, a structure or a computation. There is no question of construing the concepts of computation too narrowly or too broadly. If we wish to enlarge the class of computations beyond what was determined by these original theoretical formulations, it must not be at the expense of precision in definition.

Physical Computation

This section begins with a summary of the conditions which a physical event or object must meet in order to constitute a part of a computation. Then a set of formal definitions is given based on

Turing's original treatment,[17] and then there is a review of the physical conditions noting that the formal definitions fail if the physical conditions are not met.

The structures a computer manipulates are either atomic or composite, and the acts a computer commits, its computations, are themselves either primitive or composite. We briefly sketch the conditions on each of these four classes of entities.

All computational structures must ultimately be built out of objects which have no internal structure as far as computation is concerned. These *irreducibles* are computationally invariant, individual and indecomposable. There can be only a finite number of them, but it must be possible to generate an unlimited number of occurrences of any one. These occurrences must yield to inspection, identification, generation and annihilation by the computer without affecting the computer except as specified in an algorithm, and inspection and identification must proceed without affecting the occurrence. These occurrences may appear at any time and place suitable to the computer. And both the objects and their occurrences must be uninterpreted. They must occupy a strictly bounded region of space and time which does not change in any relevant way during the computation except as directed by the computation. By "strictly bounded" I mean that there is a fixed and constant bound on the amount of space and time needed for this storage. No matter which irreducible is stored, these bounds do not change. Any essential geometric properties of an occurrence, such as shape or orientation, must remain fixed throughout the computation, and occurrences of irreducibles must be independent of any particular energetic or temporal factors involved in physical realization. A liter of water at room temperature is an unlikely object to compute with. If it vaporizes, we lose control of its extent altogether. It is difficult to imagine how a computer could deal with the spatially invariant properties of a liter of water vapor freely dispersed. And geometric distortions, for example such as might turn an occurrence of "G" into an occurrence of "Q," must also be excluded. From the fact that the set of irreducibles must be finite in number it follows that they cannot be a dense set, but it is important to appreciate in more detail that a class of irreducibles may not have the denseness property. The distinction is between "discreteness" and "denseness" and depends upon an ordering of the set of irreducibles. The non-negative integers, although infinite in number, are a discrete set because between any two consecutive numbers there are no others. But between any two distinct rational numbers there is always another rational number different from each. Thus there can be processes of mid-point construction that are endless. Denseness must be explicitly excluded from computational formalisms. It is inconceivable that we can represent any of an infinite number of arbitrarily close objects in such a way that a computer can generate, annihilate, inspect and identify these representations with finite and strictly bounded resources.

A composite computational structure differs from an irreducible in that it has constituents. But it can have only finitely many immediate constituents and this immediate constituency relation must be completely definite in character. If composite forms are to be well-defined, then the immediate constituents of a composite must be decisively distinguishable one from another. And the transitive closure of this constituency relation must provide the conceptual means for a formal counterpart of the idea that one thing "occurs in" another. Composites must be "grounded," in that any progression starting with a composite and stepping only through immediate constituents must always terminate after a finite number of steps with an irreducible. Composites must yield always and decisively to inspection, identification of parts, replication, construction, comparison and disassembly, all by strictly bounded physical means, without affecting the computation in any other way. Arbitrary increase in degree of composition must not in principle alter any computational characteristic of the composite. And composition itself must be expressible in terms of a principle of well-formation, an idea which seems to appeal to the concept of computation for its justification. Finally, composites must be formable with the greatest possible freedom consistent with the stated conditions on well-formation.

In general, the class of computational acts may be divided into a class of operations and a class of judgments. Thus there will be primitive operations, primitive judgments, composite operations and composite judgments. The effects of a primitive act must be definite and predictable in all cases. A primitive act must have a definite beginning and a definite end and the beginning and end must be recognizable by other parts of the computational machinery. A primitive judgment must be able to determine whether a specified object is a computational structure and, if it is an irreducible, there must be a single judgment which can identify it uniquely. There must be primitive judgments to compare irreducibles for sameness or difference. There must be primitive operations to generate new occurrences of irreducibles. There must be primitive acts to discriminate among composite structures, to decompose composites into their immediate constituents and to assemble new occurrences of composites out of given occurrences. In all cases, the resources needed to initiate and conclude a primitive act must be strictly bounded. There must be a fixed and constant bound on the resources needed for the committing of the primitive act. No matter what arguments are supplied, and no matter when or where the act is committed, its demands on resources never exceed the strict bound. The effect of a primitive act must be independent of the time and position at which it occurs, provided it is admissible at that time and position. Nor may the effect of an act depend on the length of time taken to conclude it, except that each act must be allotted a non-vanishing interval of time for its completion. The resources needed to conclude a primitive act must be independent of the particular structures to which the act is applied, and therefore must be independent of any interpretation that may be placed upon these structures, whether they are irreducible or composite.

Composite acts are specified with reference to occurrences of primitive acts. In general, composition is specified by restricting the order in which primitive acts may be committed. Particular patterns of restrictions on the order of primitive acts are called "control structures." A further restriction is that no primitive act may be initiated until its immediate predecessor in the order has concluded. Such a computer is said to be "synchronous" and one which can initiate an act before its predecessor is concluded is said to be "asynchronous."

Since the principal concern here is with the physical aspects of computing, I have chosen to describe one particular formulation—the Turing machine. There are other formulations of computability theory. The lambda calculus[4] deals with functional abstraction and application of functions to arguments. The equation calculus[8] deals with definition by induction and its extensions. Post's combinatorial systems[12], Markov algorithms[10] and Chomsky's phrase-structure systems[11] all deal with the manipulation of strings. In each case, specification of a machine to realize the required processes leads to exactly the same physical constraints as does the Turing machine treatment. The interpretations that motivate these formulations are both interesting and important, but I will not consider them any further here.

A Turing machine is a physical device which acts upon a strip of tape divided uniformly into adjacent squares. A square may be blank or it may contain an occurrence of some individual symbol chosen from a fixed and finite set of symbols, called the machine's "alphabet." The machine has a tape sensor which in operation attends to exactly one square on tape. The machine can commit certain primitive acts upon this tape. It can shift the tape one square at a time past its sensor in either direction, it can record an occurrence of one of the symbols on the attended square and it can erase what is already on that square. It can determine if the attended square is or is not blank, and if not blank, it can determine which symbol occurs. A tape configuration is a sequence of marked or blank squares which includes all the marked squares on the tape.

A Turing machine embodies exactly one complex act which can be committed on any of a variety of tape configurations. The act embodied in a Turing machine is realized as a physical event, determined in part by the design of the machine and in part by what occurs initially on tape. A Turing machine can produce different results systematically only by examining and responding to the symbols occurring on its tape in a step-by-step way. In general, the decision to commit a particular act, primitive or complex, is contingent upon the occurrence of a specified configuration of symbols on tape. A general means for expressing arbitrary conditioned complex acts of this kind is achieved in a finite set of commands, each represented by a list (A, B, C, D). The first item, A, is a positive integer identifying the command. B specifies the symbol whose occurrence on the attended square must be recognized. C specifies the act to be committed and D identifies the next command. A command whose fourth member is 0 is a stop command. A set of commands is admissible if and only if no two distinct commands have the same first and second members. Finally, a particular command number is distinguished which identifies a "first" command, the beginning of the algorithm. An admissible set of commands together with an initial state completely determines an algorithm and thereby a Turing machine.

A Turing machine in action is attending to two things—the next command number and the square under its sensor. It advances its activity one step by comparing these two items with the first two members of individual commands, and a successful match determines a primitive act and a next command number. The primitive act is committed and the command number sensor is made to attend to the next command number. Persistent iteration of this stepwise advancement produces a computation.

A particular admissible set of commands completely determines a Turing machine, but further conventions are normally associated with the systematic behavior of all Turing machines. Computers are social tools, and anyone using them must know the social rules. There is first an operational convention designed to prevent the machine from running out of tape before its computation is concluded. We assume that whenever the machine is about to shift the end of the tape past its sensor it suspends operation and emits a signal which is not terminated until more tape is provided. Further conventions make the machine's behavior uniform in all situations. A strip of tape initially has only a finite number of marked squares, and a left-right tape direction is coordinated with the machine's tape-shifting capability. The left-most marked square relative to this coordination is then positioned under the tape sensor, so that the machine is attending the "first" mark on tape. The other attention aspect of the machine is directed at the initial command number and action is initiated. After the passage of some time the machine stops (perhaps). The result of the computation is the sequence of symbol occurrences whose first is found in the square under the sensor.

Even more conventions are needed if a machine is to commit familiar acts. These conventions are expressed in a systematic scheme by which entities are represented as sequences of symbol occurrences. We bundle up these conventions into a function *rep* such that *rep*(x) is a systematic tape representation for x. In general *rep* cannot be specified without knowing how the commands are to be organized, and the commands are not determined until the details of *rep* are given.

A Turing machine "applied" to a particular *rep*(x) may terminate with its sensor attending to the start of some *rep*(y). If a Turing machine T initiated with its sensor attending the first symbol of A terminates attending the first symbol of B, we write "T:A –> B." One of the most important definitions in computability theory can now be stated:

> If f is a function whose arguments are in X and whose values are in Y, then f is *Turing computable* if and only if there is a Turing machine T and a representation function *rep* such that for any x in X and y in Y, f(x) = y if and only if T:*rep*(x) –> *rep*(y).

Special attention is called to the fact that the action capabilities of Turing machines are described without reference to the conventions for the systematic use of the machines. We also note that reliance on the results of a Turing machine's computations is measured not only with regard to whether the machine acts physically as specified. Turing computations are determined as well by the simultaneous specification of an admissible set of commands and of a function *rep*. These of course are human activities having both individual and social characteristics. When we say that a certain function is computable, we have inevitably accepted a certain amount of interpretation.

A state of a Turing computation is any complete description of the tape configuration, the tape square under attention and the next command number, together with the admissible set of commands. And a Turing computation is a finite sequence of states such that the admissible set of commands justifies the machine's transition from any state in the computation to its immediate successor.

It is recognized that an admissible set of commands may itself be expressed as a sequence of symbol occurrences. These sequences may be written on tape and the tape presented to some Turing machine. In fact, a *universal* Turing machine can be specified, one whose initial tape configuration contains a representation of an admissible set of commands determining a Turing machine T, and a representation, *rep*(x), prepared for T. The universal machine, UT, then duplicates the results obtained by applying T to *rep*(x). That is,

$$\mathrm{UT}:(rep(\mathrm{T}), rep(\mathrm{x})) \;{-}{>}\; rep(\mathrm{y}) \text{ if and only if } \mathrm{T}:rep(\mathrm{x}) \;{-}{>}\; rep(\mathrm{y}).$$

In short, all the facilities which a Turing machine must have in order to scrutinize its commands are already included among the facilities it must have in order to act upon its tape in response to these commands.

The fact that a universal Turing machine can be specified is a significant fact about Turing machines and about computation in general. It means that there is a finite specification in one admissible set of commands which can direct the elaboration of any Turing computation whatever. A universal, or general purpose, Turing machine is a "most powerful" or "maximal" computer. The range of activity of a general purpose machine is as wide as the class of all possible computations. (This is Church's thesis.[3] No one has ever identified an intuitively recognizable computation that lies beyond the power of any Turing machine. But we still have no guarantee that the thesis is a law.)

The notion of a universal machine also provides a standard. A set of primitives and a means for organizing them into composites is adequate for all computations if and only if a universal machine can be specified by their means. Finally, the existence of a universal machine demonstrates that it is an intrinsic feature of algorithms that they are representable in the same way as are the ordinary structures of computation. The capability which a system has when a universal machine can be realized within it is a technical notion that corresponds to the intuitive notion that there is no limit to what the system can realize except what is found in its formal specifications. I think that this is the notion linguists have in mind when they say that the mechanism for language is "creative."

It is easy to appreciate that the physical constraints listed in the beginning of this section are essential for the construction of reliable Turing machines. A Turing machine's symbols are its irreducibles, and if occurrences of these could change in any way except as directed by the Turing machine, then Turing computation could not be well-defined. In fact, a Turing machine doesn't change these occurrences at all. It detects, recognizes and discriminates among them, without affecting them, and it can generate and annihilate them. But it never changes them. Now, we may not be able to find such absolute invariance, for arbitrarily long periods of time, in the structures which biological systems may use in their computations. I propose that it is necessarily a property of a biological computing agent, if there are any, that it attributes irreducibility *only selectively* to the objects with which it computes. The irreducible elements in a physical computation are irreducible only *relative to* a particular computer. This is a fundamentally important notion in the search for biological computation. Organisms have an interest in other organisms and in organic matter in general. If these are to enter into a living thing's computational habits, then that living thing or its computing organs will have to be able to attribute irreducibility to things which may not have this property at the physical level, even in respects that concern the living thing. It will in fact be advantageous for the living thing to be able to attribute irreducibility in as unrestricted a way as possible. This attribution must be the result of the action of a free, creative and general purpose agent!

The Turing machine's irreducibles must clearly be chosen from a finite repertoire. Otherwise we should have to have a finite tape into whose squares any of an infinite number of distinct symbol occurrences might be placed. Turing computation requires that there be no finite limit to the number of occurrences of its symbols that it may produce. A Turing machine's tape may grow in length without bound in the course of a computation. Any symbol must be writeable into any square, and once written, a symbol occurrence must retain all its properties invariant for as long as the computer may require. Occurrences must be strictly bounded in spatial requirements, else we couldn't build a machine to manipulate the physical tape itself.

In any given state of a Turing computation there are only finitely many marked squares and the smallest sequence of squares including all these and the square under scan may be taken as the machine's composite structure. Its immediate constituents are its individual squares together with whatever occurs in them. Each of the conditions required of a computer's composite structures is satisfied by a Turing machine's tape configuration. The notion of well-formation for a Turing machine's tape may be interpreted either broadly or narrowly. In the broad sense, the set of well-formed tapes is simply the free semi-group under concatenation with the alphabet of symbols as its generator. Narrowly, it is the set of tapes determined by the function *rep*.

The reader may easily convince himself that the primitive acts of a Turing machine are definite and predictable, have a clear beginning and end, are arrangeable in orderly progression, can recognize occurrences of particular symbols, can generate new occurrences and annihilate old ones, and depend upon the invariance of the machine's tape configuration. A Turing machine's activities are completely independent of the place and time at which they occur, and are also independent of the length of time taken to carry them out. And temporal and energetic factors are clearly irrelevant to the Turing machine's operation.

Furthermore, a computation is a "closed" event. For a Turing machine there are no computationally interesting consequences of a computation anywhere except on its tape. And there are no influences acting upon and determining the course of the computation except those that arise from the admissible set of commands and the precisely specified way in which such sets of commands are executed. As a matter of definition, the communication interface between the machine and all of external reality is restricted to the squares on its tape. In particular, no systematically varying global factors may influence the effects of primitive acts in any computationally interesting way. A computation is an isolated event. People are accustomed to place interpretations on the results of computations, and they are accustomed to attaching other devices to the peripheral equipment surrounding a computer. In the latter case, it is as though we attached a gadget to a particular square on the tape in such a way that each time the machine writes a 0 into that square the gadget in turn commits its act upon external reality. Such an event is, of course, a consequence of a computation, but it is not itself necessarily a computation. As far as the computer is concerned, the gadget needs no other property but this: it must embody a freely initiatable and otherwise orderly event. Alternatively, we may arrange matters so that when a Turing machine reaches a determined point in its activities, it suspends its own activity, "transfers control" to some external agent, one which may even possibly act upon the computer itself, and waits until control is transferred back to it. This is a notion of "relative" computability. It is an interesting notion for natural or biological computation, particularly in nervous systems, but I will not consider it further here.

Is the Brain a Computer?

If computation is to be implicated in any aspect of mental life, we look naturally to the brain for evidence of computational entities. In this section I will indicate that neural processes may not be rigorously interpreted as having computational characteristics. I will assume that neural activity consists of trans-membrane currents arising out of systematic changes in ionic concentrations and other subthreshold processes. Some of these are electrical and others are biochemical. The reader may find these matters discussed in Katz,[7] Shepherd[16] and Kandel.[6] There are of course many aspects of neural activity which cannot be treated here. These include special properties of the afferent systems, spinal mechanisms, the myo-neural junctions, biochemical processes underlying neural behavior, lateral inhibition, feedback loops, metabolic regulation, comparison with the immune system, studies of brain correlates with learning and behavior, ontogenesis, lateralization, and other important and relevant topics. I will refer principally to the "firing" event itself, and to the most important physical manifestations of it, together with the way in which one set of firing events induces another.

The human brain consists of about 1500cc of living tissue, shock mounted and isolated in a bath of cerebrospinal fluid, and protected by a strong bony shell. Day and night, through sleep and waking, this tissue is never quiet, and is never the same from moment to moment. Nor does it ever return to a state it has once occupied. The evidence is decisive that the brain is implicated in consciousness, sensation, intelligence and volition, and no complete study of human behavior or its nonphysical correlates may ignore the neurons in the brain.

The centerpiece of the contemporary picture of neural activity is the impulse, induced by thousands of microscopic excitations and inhibitions, distributed freely but specifically over the membrane of the cell body and its dendritic system. These local currents wax and wane exponentially, and their effects spread according to an exponential and, hence, smoothly graded, functional relationship. Voltages and ionic concentration gradients vary smoothly and systematically. The neuron's threshold responds to sustained excitation by gradually increasing, and is even sensitive to the rate of excitation. "A current of slowly rising strength may be imperceptible and not set up an impulse, even though it may rise gradually to an intensity many times greater than that at which a quickly rising square pulse is effective."[7] Because of this sensitivity to rate of excitation, "sinusoidal alternating current of very low frequency is ineffective because the rate of change of the current intensity is too low. An alternating current of very high frequency is also ineffective because its half cycles are too brief to displace the membrane potential (the small effect of each half cycle tends to be cancelled by the next half cycle during which current flows in the opposite direction)."[7] The response of a neuron is thus not only determined by extraordinarily complicated graded events, but is even sensitive to the first derivative of those events. At a given point on the membrane an excitation falls exponentially with time and spreads spatially with exponential decrement. Summation of the effects of excitation at various points takes place at the axon hillock where impulse initiation is nonlinear. But the supporting summation itself is adequately described, as far as we know, as a graded process.

Neither can we say that the impulse is the only significant event, and that subthreshold activity is merely its prelude. Perkel and Bullock write that ". . . the role of impulses in the central nervous system in representing and transforming information has seldom been established and is nowhere investigated to a satisfactory degree of completeness. Moreover, several kinds of evidence, although somewhat indirect, point strongly to the importance of other, nonimpulse vehicles for carrying information in the brain according to their corresponding coding schemes; the importance of such nonimpulse codes may well surpass that of the 'classical' nerve impulse."[11] "The suggestion has been made and must be entertained seriously that it is the impulses that are best regarded as the epiphenomena (at least in some parts of the central nervous system) and that only through understanding the properties and interactions of the electrical waves with the anatomical substrate will we arrive at a satisfactory understanding of the higher behavioral and mental processes."[11]

It is experimentally observed that the threshold and the refractory period determine a range of rates or frequencies at which the neuron can discharge impulses. The observable fact is that these rates are selected from a dense set. It is also clear that these rates may be expected to vary incrementally with graded changes in activity at different synaptic sites on the neural membrane. Complex patterns of such graded stimuli may produce complex patterns of variation in firing rates.

In their survey of neural codes, Perkel and Bullock review those classes of neural events that have been observed in the nervous system, or whose occurrence may reasonably be deduced from what is observed, that may constitute neural codes, or may be essentially implicated in the generation, transmission or transformation of coded signals. The labeled line, which identifies the quality and location of the signal source, and the onset, or unit timing code, where the first impulse is the bearer of information, are the simplest. Principal attention is given, however, to the rate or frequency code and I will try to summarize some of their observations. A particular occurrence of a firing rate may be determined by a single pulse interval or by an average taken over an interval of time. The physical events needed to realize either of these determinations can hardly be the primitive acts of a computing machine.

There may be a background frequency, the increment being the carrier of information, or else the absolute frequency itself may constitute the signal. Rates may vary slowly or rapidly. For a mean rate code, information may be temporally weighted, more recent pulse intervals contributing more to the coded signal. In yet another situation, "both the mean rate and the distribution of intervals in the postsynaptic cell were in some circumstances highly sensitive to the standard deviation of impulse intervals in the synaptic input. This sensitivity was combined with a sensitivity to mean input rate."[11]

The temporal coding of micropatterns without changing mean pulse rate has also been observed to be significant. Certain muscle contractions may be "many times greater for a train of alternately long and short intervals than for one of uniform intervals at the same mean rate."[10] These rate sensitive events are not compatible with a computational interpretation, since they are all continuously varying. In addition, the end result of the action of those neural assemblies implicated, say, with motion, is to control actions that are intrinsically graded. It would indeed be a spectacular situation if these graded phenomena could be controlled and apprehended by a mechanism which was denied any understanding of the concept of continuity.

Perkel and Bullock point out that microstructure codes require extremely precise timing, and that pattern codes may therefore not be so widely used. We may observe that human activities sometimes require extraordinary precision in timing which must originate in the higher brain centers. In addition, the representation, discrimination and manipulation of psychologically significant spatio-temporal events in the brain require a representation of these dimensions, and the nervous system has only anatomical specificity and temporal coding available to it to achieve this spatio-temporal representation. Neocortex, for instance, is a two-dimensional sheet, and so therefore are the motor and somatosensory homunculi that have been identified.

In parallel fibers coordination of impulse patterns becomes an important issue. "The use of systematic delays to encode the desired time relations of the effector action . . . is probably a general principle of parallel line or ensemble coordination."[11] If the coordination of distinct timing patterns to an arbitrarily fine degree is found to be of significance, we shall have another instance of continuously varying signals arising in the brain.

In almost every respect it is difficult to accept the proposition that the brain is a computer. The capacity of the neuron to fire at arbitrary rates provides elementary units whose occurrences do not constitute a discrete set, let alone a finite set discretely organized. The microstructure of firing patterns is observed to constitute a rich variety of events, and shows considerable structure and regularity.

In addition, neural accommodation and the biochemical events that are believed to subserve learning, make it difficult to believe that neurons have physical invariants of the kind needed to constitute computations. And the difficulties with localization theories make the computational requirement for closure very hard to satisfy in any particular aspect of brain function. The arousal systems in the brain stem and mid-brain reticular formations appear to be so widely connected in the thalamus and cortex that it is difficult to imagine how arbitrarily remote global factors may be excluded from important aspects of brain function.

More broadly, if we try to justify any computational interpretation of neural function, we shall have to explain their precursors in simpler organisms, and show how these computational characteristics are transmitted in ontogeny and phylogeny. Otherwise, we have to understand that computational habits are injected ad hoc. We may say that computation is an "emergent," but this of course has no explanatory value whatever.

We may say that the brain is really a computer whose irreducibles are rate codes. After all, it may be claimed, neurons count integral numbers of spikes. Nothing in this picture is justified. We do not know that neurons, or neural assemblies, in general "count" spikes or anything else. So little is known about the integration of complex neural activity, apart from the fact that it is more often than not graded, that there is no experimental ground to reject the position that firing rates are determining factors at arbitrary levels of neural integration. This observation is reinforced by Perkel and Bullock's statement that "no limit can be set, at least on the dimension that makes important conscious decisions 'high,' to the upper level that single-unit thresholds may be responsible for."[11] If single-unit thresholds can contribute to arbitrarily abstract levels of integration, then so indeed may rate codes and their microstructure, and there is then no basis for assuming that the contribution of membrane activity to brain processes is computational. Those discrete structures and functions that have been identified may do no more than isolate physically the continuously varying influences that are required for reliable, predictable and repeatable integration into more complex continuously varying processes. Presently available evidence, in short, fails decisively to support the conclusion that the brain is a computer in a well defined physical sense, or that there is any significant physical aspect of brain function that is computational in character.

One may consider that the brain is in fact a device which operates on graded domains in such a way that the results either physically or in consciousness have the form of limits of convergent processes. The brain would be a graded processor that "converges" to a computer. We would then not expect to find computational activity in the brain. Rather, we would look for convergent processes. I cannot consider this alternative any further here except to observe that a theoretical framework for this proposal may be found in the recent studies of Scott[14,15] and others, studies which formulate the basic concepts of computation in terms of approximations, continuity and other topological notions.

Computation and Attention

I will be speaking now about my own personal experience with awareness. I distinguish awareness from attending. Awareness is an experience and attending is an act. One can be aware of the object of an act of attention, but one can also be aware without attending. One cannot, however, attend without being aware. Attention is thus conscious attention.

The central aspect of awareness is the unitary character of the form that attending takes. A particular occasion of attending will normally involve a variety of objects, aspects of these objects, relations among them and so on. But there is a threshold above which some single entity—an object, a feeling, an aspect, for example—is apprehended in a unitary way. In Walker's words[18] consciousness is an onset phenomenon. Below this threshold we cannot truly be said to be attending, although we may well be aware. The object of attention may have constituent structure that I am aware of, but the form the entity takes in my attention has a unitary character that it retains as long as it remains the focus of attention, a character that gives it a prominence in my awareness that no other object simultaneously present to awareness can have. My attention may shift to one of the constituents of that entity, but then that constituent itself becomes the focus of attention. My attention may shift to the fact that such and such is indeed a constituent. Then that particular constituency relation will assume the unitary focal character.

I may attend to an event that originates in my periphery—my thumb, the position of a limb, my posture, a soothing sound, a bright light. I may attend to yesterday, to tomorrow, to the future, the past, now (with some difficulty), this, that, here, there, me, him, them, a piece of reasoning, a drink of water, an impulse, an image, red, color, an impression, impressions, my dog, dogs. In particular, my attention can fixate on one thing through an interval of time, a peculiarly difficult notion for those who want to find attention in the physical structures in the brain, since we have seen that at least at the level of membrane processes there are no readily detectable invariants. It is then difficult to suppose that these membrane events support the invariants needed to account for our direct experience with attention.

My attending to an entity can become more intensely fixated, and as this intensity increases the entity dominates my awareness more completely. I can attend to an arbitrarily selected point on a homogeneous surface, to a pair of points, to a set of points, to systematic arrays of points, to the idea of a random array of points, (but not perhaps to a particular such random array). I can attend to lines and curves connecting points on this surface, to neighborhoods containing points, to all the points on the surface, and to "no points on the surface." But no matter how diffuse the background details, the form taken by the foreground object of my focal awareness is always unitary, unique and invariant.

The physical shape of the object of awareness may be smoothly graded, as when I contemplate any moving object. But the form of attending is not the same as the substance of what is attended to. This form is an absolute that endures through time, as when I continue to attend to the motion of an object.

I notice that my awareness changes easily. Under suitable excitation my attention transfers abruptly from the cat's fur to the ringing telephone. With guidance it transfers from cat's fur to cat's paw to cat's claw to skin scratch to iodine bottle. Just as easily it shifts from integers to rational numbers to real numbers to vector spaces. But the form of these changes is always abrupt and discontinuous. In shifting from one entity to another my attention does not pass through a graded sequence of stages of attention, each differing but little from its neighbors. A shift in attention is the kind of act a computing agent might commit.

I am on occasion more or less aware of the attending aspect itself of the act of attention, but this aspect does not thrive as I attend more sharply to the object of attention. On the other hand, I may try to attend to the act T of attending. I can easily do this, but the object of my attention then ceases to be the object of T and comes to be T itself. My new state of awareness has as its object the traces of T. Curiously, I can iterate this shifting of attention to the act of attention as many times as I like, with little effect on my attention. There are aspects of the description of this phenomenon that remain unchanged as I move from one level (of attending to attending) to another level (of attending to attending), aspects that may be expressed in the computational idea of recursion.

There seems to be a substrate out of which my consciousness fetches the focus of attention and its supporting structures. I can feel this substrate, I am aware of it, I know it's there. I am not permitted, however, to attend to its ingredients or its details. It is definitely graded in character. Perhaps, this is why I cannot attend to any differentiated part of it—there is no basis for differentiation. Some unknown process extracts from this substrate the discrete structures that become the objects of my attention and gives them their computational character. The focal entity in the foreground of my attending and its immediate neighbors are individuated, and support each other's definiteness. The stream of awareness is organized in its background as a system of graded processes and relationships. Pribram[13] has made significant suggestions about a possible continuously varying structure which may be related to this background substrate.

Abrupt and dramatic changes in attention may be brought about by discontinuities in the stream of data coming along the afferent pathways, or they may arise as a result of internal disturbance. In such cases, one would not quarrel with the judgment that the changes in the state of attending are discrete. In the more usual situation, temporally consecutive but distinct stages of attending are alike, on the average, except for some small aspect. Such shifts in attention may consist in change from the attended object to another object related to it, a separation of the object into parts or properties, a shift by generalization, a shift to a relation, etc. Each such change seems to me to occur abruptly as far as attention is concerned. Successive stages in awareness are similar to the successive stages in a computation, as far as their foreground structure is concerned. If the background of awareness flows continuously, the foreground lurches forward in saltatory leaps, producing state changes which I can only become aware of by a suppression of the background and a sharpening of my attending to what is in the foreground.

These observations suggest that there is discoverable order and regularity in attending and in the processes of changing attention. It is proposed here that structures and processes involved in the form of attention and in attention changes are computational—systematic manipulations of pure form.

Indeed, natural computation at an essentially abstract level has already been sought and found. A generative transformational grammar for a natural language is a certain kind of computational specification.[2] It assumes that the production and understanding of language involve the recognition and manipulation of structures having definite form—finite and discrete—and that there are certain kinds of rules that represent relations among these structures. The word "generative" means that these structures, rules and relations are computational, and that the effect of applying rules to structures is managed, at least in part, by computational control mechanisms. The notion of a generative transformational grammar may not survive very far into the future, but at least two aspects of it are of considerable importance. One is that our experience with language may naturally be represented in terms of a finite system of categories which not only classify linguistic objects, but also reflect both their assembly into complex objects and their susceptibility to interpretation. The other is that occurrences of objects belonging to these categories can be organized in such a way that linguistic objects express simultaneously the invariance of certain categorial relations and the free variation of others. "I gave him some soup" and "What did I give him?" have certain meaning relations in common, although syntactic functions are quite distinct. The matter does not end with language. Lenneberg[9] has suggested that "All vertebrates are equipped to superimpose categories of functional equivalence upon stimulus configurations, to classify objects in such a way that a single type of response is given to any one member of a particular stimulus category." He observes that most higher animals have a species-specific and useful capacity for discrimination, and that "Most primates and probably many species in other mammalian orders have the capacity to relate various categories to one another and thus to respond to relations between things rather than to things themselves . . ."[9] Categorization, differentiation and transformation become basic organizing principles for higher animals, and perhaps for living matter in general. The membrane of a neuron, for example, differentiates ions into species according to their potential effects on the membrane, and makes use of this categorization to transform local aspects of its configuration into other configurations with "adaptive value" for the neuron.

It is proposed here that the mechanism that administers attention organizes its interaction with the subject matter in discrete, finite and unitary ways. And shifts in attention are formally specified as preserving certain relations among these categories, while other aspects of the whole situation are allowed to change. Like cognitive differentiation and language, the pure uninterpreted form of attention is computational, while its substance, its interpretation, is non-computable.

I also notice that there seems to be great freedom in my ability to select objects of focal awareness. It seems that causal laws do not operate at the level at which the attending mechanism manipulates pure form. I may apprehend a great deal of causality in the substance of what I attend to. But the way in which factors outside attention control that attention is as mysterious as the computational form of that attention.

I think it is also true that the discretized and computational structures determining the form of my attending are terminal, or "dead end" in character. The form of what is attended supports a representation, a picture or image. An act of attending is in this respect like an act of measurement. The form of the result is an invariant which constitutes a systematic description to the attender of some specific aspect of what is attended. Once obtained, the form of this description does not change. A new description may appear, sufficiently close to its predecessor to give the illusion of a changing description, but each consecutive description is static and fixed. I attend to the motion of a particle and am aware that it is accelerating. My attending may have been fixed on the particle at one time, but if I know it to be accelerating, then my attending has shifted to its acceleration. From the act of measuring its position I have moved to the act of measuring its momentum. I do not do this in graded stages. It is the form of attending that determines that a measurement will have a definite, unitary and invariant value.

Discussion

In the earlier pages of this essay I took pains to elaborate in considerable detail exactly what it means for a physical process to be a computation, and I then observed that at the level of the nerve impulse nervous tissue cannot be understood to act computationally. Finally, I described my own experience with attention and suggested that when stripped of all interpretation, the forms of the attention mechanism appear to be computational. If, in fact, living systems fail in general to reveal computing mechanisms in their organization, then this fact may have significant consequences for the study of mind.

It is clear from the simple experience of attending to motion or to any other smoothly graded phenomenon that the substance of attention is not in general computable, although its form seems to be describable in computational terms. We may interpret these facts as indicating that computational structures arise exclusively at the interface at which reality manifests itself to awareness, that these computable forms are the instruments of attending, observing and measuring and nothing more. The anatomist's notion of articulation comes to mind. Skeletal structures articulate certain biological forms on which flesh and blood will later be hung. Computable structures articulate certain more general forms, on which reality will eventually be hung.

If these observations are correct, they may have consequences for experimental parapsychology as well. If we postulate a level of organization of reality that is computational and non-physical, we may try to refer paranormal phenomena to this level. For example, it is possible that paranormal communication occurs most easily in connection with the emergence of an act of conscious attention, in the state in which what we are about to attend to is arising from the unattended substrate under the action of forces presently unknown and by processes only faintly perceived. This act of arousing a unitary and attendable object out of the unconscious substrate is similar to the physical process of making a measurement. It may be an important target for parapsychological experiment within the framework proposed here. Griest[5] has already reported that language may affect psi functioning. We do not know whether this effect is due to the substance of language or to its form or grammar.

In addition, computing devices are widely used in experimental parapsychology, where the nature of the experimenter–subject–apparatus interactions is ill-understood or not understood at all. If computational processes are of fundamental and universal significance in the relation of consciousness to reality, then that system of ill-understood interactions may need to be expanded to include the computer, not merely as a piece of apparatus, but as a physical approximation to an abstract mechanism for which the mind may have an absolutely unique affinity. This latter possibility may itself be subject to experimental investigation.

BIBLIOGRAPHY

1. Chomsky, N., "Three models for the description of language," *I.R.E. Trans. on Infor. Theory,* Vol. *IT-2* (1956), pp. 113–124.
2. Chomsky, N., *Aspects of the Theory of Syntax,* M.I.T. Press (1965).

3. Church, A., "An unsolvable problem of elementary number theory." *Amer. Journal of Math.*, vol. *58* (1936), pp. 345–363.
4. Church, A., *The Calculi of Lambda Conversion*, Annals of Mathematics Studies No. 6, Princeton Univ. Press (1941).
5. Griest, W., "Psi speech communication and cognition," in *Extrasensory Ecology: Parapsychology and Anthropology*, edited by J. K. Long, Scarecrow Press (1977).
6. Kandel, E. R., *Cellular Basis of Behavior*, W. H. Freeman (1976).
7. Katz, B., *Nerve, Muscle, and Synapse*, McGraw Hill (1966).
8. Kleene, S. C., "General recursive functions of natural numbers." *Math. Ann.*, vol. *112* (1936), pp. 727–742.
9. Lenneberg, E. H., *The Biological Foundations of Language*, Wiley (1967).
10. Markov, A. A., *Theory of Algorithms*, Translated by the Office of Technical Services, U.S. Dept. of Commerce (1952).
11. Perkel, D. H. and T. H. Bullock, "Neural coding," in *Neurosciences Research Symposium Summaries*, edited by F. O. Schmitt, T. Melnechuk, G. C. Quarton and G. Adelman, M. I. T. Press (1969).
12. Post, E., "Formal reductions of the general combinatorial decision problem," *Amer. Journal of Math.*, vol. *65* (1943), pp. 197–215.
13. Pribram, K. H., *Languages of the Brain*, Prentice-Hall (1971).
14. Scott, D., "Outline of a mathematical theory of computation," *Proc. Fourth Annual Princeton Conference on Info. Sciences and Systems*, Princeton Univ. (1970), pp. 169–176.
15. Scott, D., "Models for various type-free calculi," in *Logic, Methodology and Philosophy of Science IV*, edited by Suppes, P., L. Henkin, A. Joja and G. C. Moisil, North-Holland (1973), pp. 157–187.
16. Shepherd, G. M., *The Synaptic Organization of the Brain*, Oxford Univ. Press (1974).
17. Turing, A. M., "On computable numbers with an application to the entscheidungsproblem," *Proc. London Math. Soc.*, ser. 2, vol. *42* (1937), pp. 230–266. Correction in vol. 43, pp. 544ff.
18. Walker, E. H., "The nature of consciousness," *Mathematical Biosciences*, vol. 7 (1970), pp. 131–178.

DISCUSSION

Morris: I would like to hear you comment more on your notion of computation as a physical event. I'm not challenging it. I just would like to hear a little bit more about what's persuaded you over the twenty-five years that it must be a physical event.

Storm: I think it's related to my problem with causality. I have a problem with causality in general and it's fixed, particularly in respect to computation. Computations, as we ordinarily see them, involve submitting one set of structures to a computer; then we run down to the ready room the next day and we get the printout, and we get something very different than we got before. It is, of course, possible in principle, that we could say that we just have pairs of structures and they're computationally related. But the concept of computation for me is the concept of an event—that I have something in my hands today, a column of numbers. We'll take a very simple example because its paradigmatic. I have a column of numbers, and that's one form and I submit that to a computing machine and I get back the sum. Well, something must happen for me to get back the sum and I don't know how causality figures in determining the sum—it's as though I were making a measurement, I were submitting some numbers and having the computer do something like that for me. I'm saying, "Here are the numbers, and now manipulate, rearrange, reorganize," and the machine comes back and gives me something different and I see that it takes time to do that. It takes expenditure of energy to rearrange things in that way. Rearrangement of form, the spatial and temporal redistribution of form even when it's uninterpreted, takes time and it takes the expenditure of energy and therefore, the expenditure of time. That's what I'm trying to say.

MORRIS: Yes, this is what I thought it might boil down to, that computation must take place in space time. Is that what you're trying to say?

STORM: Yes.

BYERS: Is it possible to have multiple, simultaneous attending at the same level of brain functioning?

STORM: Yes it is. As far as I'm concerned, in my personal experience, it's a problem. I want to use the word "hysteria" in connection with that phenomenon. I don't like that phenomenon when it happens. The next thing that happens when I find myself attending, truly attending simultaneously to more than one thing, I start feeling as if I don't have it all together any more—the question perhaps of parallel processing. I don't know whether the purely abstract apparatus I'm talking about is able to process in parallel. But I know I'm happiest when I feel most together, when it's one thing that I'm dealing with; that my attention is on one thing. Now that's the question of simultaneity. The question of levels is a question of whether I am aware at this level. I'm aware at the present moment. I'm attending to the fact that words are coming out and I'm looking at you and in the background I'm thinking about what you just asked me. But I know that if I peer down underneath there are going to be other subordinate levels of structure down there which I could get at by attending to different things. I could start picking my ideas out, or I could start talking about the color of your coat.

UNIDENTIFIED VOICE: There was some research which I would not try to cite at the moment, but which I read a couple of years ago, about someone who attacked the problem of whether one could do two things at the same time—in this case, read and write quite different stuff, and found that initially that was not possible. That is, the literature seemed to suggest that was not possible, but he found that with training it was, and the reason that it had seemed that it wasn't was that nobody had ever taken the time to get over the hump of learning to be trained to do that.

STORM: I wonder if that's ever really going to be a matter of attending to two things at once or if we might not find it to be a matter of being able to perform one thing or perform two things and I wonder what the attention experience would be like. My friends tell me that sometimes I act as though I'm attending to several things at once, and I really wonder what it would feel like. I think I could not say that myself. My own personal experience about attending—if you said, "What does it feel like, Edward, to attend to two things simultaneously," that's one place where no matter how bold I felt, I'd have to say I don't know. My own experience is I can do two things at once. I can write out by habit my first name on the paper while I'm talking to you, but that's something different. It's purely formal attending.

MORRIS: Do you envision brain function as a discrete series of events?

STORM: I'm not a physiologist, but from what I have read, I am beginning to elaborate an opinion that we don't know anything about it. I certainly am elaborating the opinion that I am very unhappy about the idea that it's the impulse that's the important thing. I have the greatest respect for neurophysiologists. I think they've got one of the most incredibly difficult physical science tasks of all time, and hats off to everything they try. I think they should be given boundless amounts of money to do it, but I think lots of them are on the wrong track in rigidly fixing their attention on this discrete impulse. I think it's maybe an overstatement for one of those fellows to say that its's an epiphenomenon. I don't think that's reasonable, because a lot of experiments on innervation of muscle tissue and other matters like that, where onset is certainly an important thing, show that an individual neuron's single firing is enough to trigger a wing-beat in an insect, and so you can't say it's an epiphenomenon. But I'm not talking about the operational issue. I'm talking about my mental life and I think at the level of mental life, it might be the case that all the discretizing and digitalizing in the brain is epiphenomenal, that its important feature is that it reflects a continuous and graded base out of which a purely abstract mechanism of some completely unknown kind at the present time fetches up what I'm attending to.

Nash: You have divided consciousness into a series of sensations with a background of awareness. According to Buddhist philosophy, there is only a series of sensations, and the background of continuous awareness is an illusion.

Storm: I think the only fair way I can answer it is to tell you what I personally think and feel, and it is the feelings that matter to me. I would also go this far—this is probably the biggest thing that has attracted me to try to make friends with parapsychologists, and that is it is not the discrete finite and bounded and therefore rational and logical structures that make me feel good and happy and make me resonate with reality and make me get it together. We could determine it on the basis of individual experience, and at the moment, I have no experience with Buddhism. But as far as my own personal experience goes, I think that there is an element of alienation in the discretizing. It can't be foregone. I think I'd be in trouble if I couldn't compute in my mind, but I also think that I get in big trouble if I don't do anything but compute. And I think there are social, educationl, political and cultural consequences of this observation as well. There is a theorem in recursive function theory that intrinsically relates the concepts of deductive logic to computation. If a set of sentences in a form of logic is universally valid, that is, if they all come out true under all interpretations, then there is a Turing machine that will determine that fact in a finite number of steps. That's an absolutely shattering result, I think. For every finite set of sentences which is logically valid, there is a Turing machine which will determine that validity in a finite number of steps. And so if I decide that boundaries have to be put on computation in the human personality, then that's going to mean that boundaries will have to be put on discrete, deductive logic as well.

Morris: When I was a graduate student at Duke, I was a subject in a really nasty perception experiment done by Nikolai Khokhlov. They would present to me simultaneously a tone and a colored square, and I was allowed to feel the texture of that colored square for one second. All three pieces of information came at once—tone, color and texture. I was trained along the gradients and asked to make judgments after the end of that second as to which level of stimulus I had just experienced for each of the three sense modalities. In other words, all three sense modalities were activated within that second, and I certainly was trying like hell to attend to all three things. The way I had to solve the problem was simply to store everything quickly and then sequentially process the information consciously afterwards. I think the relevance of this is that it seems to be possible to acquire a lot of information simultaneously and at least storing it in short term memory. At any given moment you only know you've been attending to something on the basis of what you retrieve from short term memory.

STORM: First of all, with respect to attending and retrieving and storing—on the Turing machine, there's a tiny little amount of memory in the state, but there's a lot of memory on the tape, and when we speak about memory in these regards, I really would like to know where memory is in the brain. And are we really remembering in the sense that we use that word? Does memory mean the same for us as it means for the devices we currently use?

The other remark I want to make in response to your question is that I can't understand anything about perception. I thought I was getting a handle on perception and someone loaned me a book by James Gibson. The psychologists here, I'm sure, know about James Gibson. I read the book and said, I think there's something very right about this and also there's something very wrong about this. James Gibson believes that the experience of perception has constituents. As I put my hand here—that act of perception involves implicitly both things I'm touching—the fingers, my arm, my hands and then all the way up to the brain, and then the signals that come from the brain back down to the arms through to the fingers. As we know now, every time you receive something, signals come from higher brain centers and modify the receptors, so that the information we're getting is constantly being modulated. And it's not clear to me that a good bit of processing and integration, even cross modal information isn't taking place in the peripheral system. One experiment I remember reading about was by Morrell. He had a patient with an open skull, and he found the neuron that responded to certain colors, and changed colors and the neuron responded with a different profile three times. The excitation of the neuron took on a characteristic shape for characteristic colors. And then they found another astonishing thing. If they accompanied the color presentation with sound in different places, the statistical profile of repetitions of the firing pattern changed systematically with the position of the sound source, and this was a neuron that was in the visual system.