Research Issues in Flexible Computing

Two Presentations in Japan

Bruce J. MacLennan
Computer Science Department
University of Tennessee, Knoxville

Abstract

This report contains the text of two presentations made in Japan in 1991, both of which deal with the Japanese "Real World Computing Project" (previously known as the "New Information Processing Technology," and informally as the "Sixth Generation Project").

(1) "Flexible Computing: How to Make it Succeed" (invited presentation, Institute for Supercomputing Research workshop, *New Directions in Supercomputing*): Many applications require the flexible processing of large amounts of ambiguous, incomplete, or redundant information, including images, speech and natural language. Recent advances have shown that many of these problems can be effectively solved by *emergent computation*, which is the exploitation of the self-organizing, collective and cooperative phenomena arising from the interaction of large numbers of simple computational elements obeying local dynamical laws. Accomplishing flexible computing will require basic research in three areas. *Theory:* We need to understand the dynamical and computational properties of systems with very high degrees of parallelism (more than a million elements). *Software:* We need to understand the representation and processing of subsymbolic and symbolic information in the brain. *Hardware:* We need to be able to implement systems having a million to a billion analog processors.

(2) "The Emergence of Symbolic Processes from the Subsymbolic Substrate" (panel presentation, MITI, *International Symposium on New Information Processing Technologies '91*): A central question for the success of neural network technology is the relation of symbolic processes (e.g., language and logic) to the underlying subsymbolic processes (e.g., pattern recognition, analogical reasoning and learning). This is not simply an issue of integrating neural networks with conventional expert system technology. Human symbolic cognition is flexible because it is not purely formal, and because it retains some of the "softness" of the subsymbolic processes. If we want our computers to be as flexible as people, then we need to understand the emergence of the discrete and symbolic from the continuous and subsymbolic.

1 Flexible Computing: How to Make it Succeed[1]

1.1 The New Information Processing Technology

1.1.1 Introduction

In Tokyo on March 13–14, 1991, MITI sponsored the *International Symposium on New Information Processing Technologies '91*, or *NIPT '91*.[2] At this symposium the NIPT Research Committee reported on the results of a two-year preliminary study for an R&D project to develop the computer-technology foundations of "the information network society of the 21st century" (Ishii, 1991). This technology comprises: (1) flexible computing ("intuitive information processing"), (2) adaptive computing, and (3) massively parallel computing, including optical computers. This article will attempt: (1) to explain the importance of new information processing technology, (2) to discuss its potential impact on computing, and (3) to suggest research necessary to its successful implementation.

1.1.2 Artificial Intelligence: Old and New

It is now widely acknowledged that traditional AI (artificial intelligence) has failed to live up to its promises. Although there are many reasons for this, a central one seems to be the inadequacy of rule-based representations of knowledge. The essence of the problem can be understood as follows. Anyone who has tried to write a set of administrative rules or procedures knows that there are always exceptional situations in which a rule should not be applied. Sometimes one can write additional rules to handle these exceptions, but then one can usually anticipate exceptions to these rules, and so forth. One may debate whether or not there can be an end to the exceptions to the exceptions to the exceptions..., but it is clear that an attempt to have a exhaustive set of rules will lead to so many rules that the system will be unwieldy and unusable. In actual practice, we do not try to write rules that cover every possible situation. Instead, apparent exceptions to the rules are submitted to a person or committee, who uses its judgement to decide if the rule is to be applied. Centuries of experience have shown the wisdom of having rule use tempered by human judgement.

The difficulty comes with rule-based artificial intelligence systems, since in these there is no human to whom we can appeal for final judgement. The designer of such a system is faced with a dilemma: if there are too few rules,

[1] Revision of "Flexible Computing in the 21st Century," which appears in *Vector Register*, **4**, 3 (Fall 1991), pp. 8–12. Based on an invited presentation at *New Directions for Supercomputing*, the fifth annual workshop, of the Institute for Supercomputing Research, 15F INUI Bldg. Kachidoki, 1-13-1 Kachidoki, Chuo-ku, Tokyo 104, Japan.

[2] This project has since been renamed "Real World Computing" (RWC).

then the system's behavior will be *brittle*, that is, minor exceptions or novelties will cause it to fail catastrophically (i.e., to behave stupidly). On the other hand, if the designer tries to include rules to cover all the possibilities, then the system will be very slow, since processing time tends to increase exponentially in the number of rules. Thus we find artificial intelligence systems that tax the capabilities of our best supercomputers, yet are still inadequate in many ways.

Growing recognition of the limitations of existing AI technology, and especially of rule-based knowledge representation, has led to the development of a very different approach, which we have called "the new AI" (MacLennan, 1988). Related terms, differing in emphasis and focus, are *connectionism, artificial neural networks, flexible information processing*, and *emergent computation*. The new AI is at the heart of the new information processing technology, but, to see why, it will be necessary to consider cognition in humans and other animals.

1.1.3 Human Cognition: Symbolic and Subsymbolic

One clue to the solution to brittleness can be found in the "cognitive inversion" of the old AI. By this we mean that the old AI did best what people do poorly, and did poorly those things that people, and even lower animals, do well. For example, some of the first successes in AI were with activities such as proving theorems in symbolic logic, playing games such as checkers and chess, and analyzing blood. These are all specialized skills that comparatively few people acquire. On the other hand, skills common to almost all people, such as the ability to recognize a familiar face, or to understand spoken or written language, or to find one's way through a forest, still defeat our most powerful computers. To cite a clear example, our most powerful supercomputers are insufficient to give an autonomous vehicle the same competence in finding its way through the natural environment as is exhibited by a mouse.

The lesson to be learned from this cognitive inversion is that brains process information in a very different way from conventional computers. Further evidence is provided by the "100 Step Rule" (Feldman & Ballard, 1982): If we divide the time it takes a person to perform a simple perceptual task, such as recognizing a familiar face, by the response time of a neuron, we find that there can be at most about 100 sequential steps between the stimulus and the response. A conventional computer cannot do much in 100 steps, certainly not something as complicated as recognizing a face. One reason that brains are able to accomplish this is that they use massive parallelism, but on a scale well beyond the capabilities of current supercomputers. For example, even in the first stage of visual processing we have approximately 100 million receptor cells driving one million ganglion cells, each of which is an independent, but slow-speed, low-precision analog processor.

The source of the cognitive inversion can be understood in terms of the distinction between *symbolic* and *subsymbolic cognition* (Smolensky, 1988). Symbolic cognition includes most of the faculties we think of as uniquely human, in-

cluding language use, explicit reasoning and formal symbol manipulation (e.g., arithmetic, algebra). In contrast, the other animals share our subsymbolic cognitive capacities, which include perception, sensory-motor coordination, pattern recognition, and associative memory.

Since conventional computers are in effect programmable discrete symbol manipulators, they are well suited to the implementation of cognitive processes that are completely formal and require great precision. On the other hand, such a computer must implement subsymbolic cognition in terms of the symbolic operations it provides. Thus, on conventional computers subsymbolic processes are relatively less efficient than symbolic processes. To understand the effect of this, imagine that we had to calculate by hand the procedures required to drive an automobile, that is, the exact trajectories of our hands, feet, etc. in response to incoming sights and sounds. Given the thousands or millions of calculations required and the very slow rate at which we do arithmetic by hand, we would never be able to drive at all. This is the reason for the cognitive inversion of the old AI.

On the other hand, brains are composed of very large numbers of relatively slow, low precision, simple analog processors. These are well-suited to subsymbolic cognition, such as perception, recognition, coordination, association, and pattern matching, but poorly suited to symbolic processes, such as formal logic and mathematics. It takes practice and effort to get our flexible, subsymbolic brains to emulate rigid, symbolic machines. Almost anyone can walk, tie a shoe or recognize a face; comparatively few can prove a theorem in mathematics.

One might conclude from the foregoing that although brain-style computers (*neurocomputers*) might be more suitable for subsymbolic processing, the conventional digital computer is still preferable for symbolic processing. However, such a conclusion ignores the issue of flexible rule use. We saw before that the old AI was limited by the necessity of rules to cover all contingencies, but that human rule use was always against a background of non-rule-based judgement. Now we can see why. Human symbolic processing avoids brittleness by partaking of the flexibility of the underlying subsymbolic processes in terms of which it is implemented.

These considerations have led the NIPT Research Committee and others to conclude that 21st century society will need a new, flexible information processing technology, and that this technology will require a new kind of computer, the neurocomputer, to complement the familiar digital computer.

Conventional information processing contrasts with flexible information processing in several ways:

	Conventional	Flexible
computation	narrow but deep	wide but shallow
state space	high precision digital	low precision analog
processors	complex and fast	simple and slow
representations	localized	distributed

Each kind of information processing has its characteristic strengths and weak-

Table 1: Approximate Numbers of Neurons[3]

human	$\sim 30 \times 10^9$ neurons
rat	$\sim 65 \times 10^6$ neurons
1 sq. mm. cortex	$\sim 150 \times 10^3$ neurons
	$\sim 10^9$ synapses

nesses. We argue elsewhere (MacLennan, in press-a) that the distinct advantages of conventional information processing are *definiteness, reliability* and *reproducibility*, whereas the distinct advangages of flexible information processing are *flexibility, robustness, adaptiveness* and *responsiveness*.

1.2 Research Issues

The successful implementation of the new information processing technology requires basic research in three areas, which may be loosely termed *hardware, software* and *theory*. We consider briefly some key problems.

1.2.1 Hardware

Achieving the goals of the new information processing technology will require a new kind of computer, with principles of operation more like those of the brain. Thus, instead of a comparatively small number of fast, high precision processors, future neurocomputers will make use of a very large number of possibly slow, low precision processors. For example, neurons take approximately 10 msec. to fire, and are estimated to perform analog computations with about one digit of precision. On the other hand, even an animal as simple as a bee may have a million neurons in its brain, and the human brain is estimated to comprise thirty billion neurons (Table 1, Figure 1). It is not unreasonable to suppose that if we want our neurocomputers to behave at least as intelligently as bees, then they will have to be capable of implementing neural networks of comparable size. Thus, the hardware goal for the new information processing technology should be to implement between a million and a billion neurons with at least the speed and precision of biological neurons. We put this conclusion in the form of a strong claim: *If an implementation technology will not scale up to a million to a billion processors with a billion to a trillion connections, then it is not on the trajectory to the new information processing technology.*

[3]Sources: Changeux (1985, pp. 50–51), Shepherd (1990, p. 392). An area of mammalian cortex of size $25\mu \times 30\mu$ contains approximately 110 neurons, regardless of the thickness of the cortex (though in primates the visual cortex has 2 to $2\frac{1}{2}$ times this number). We are assuming 6000–7000 synapses per neuron. DARPA (1988) suggests that a bee's brain has approximately a million neurons, but this seems inconsistent with the data we have, and DARPA (1988) cites no source.

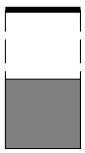

Figure 1: Neurons and synapses in 1 cm^2 of cortex. To help visualize the density of neurons and synapses in the brain, the above square represents a 1 cm^2 section of mammalian cortex, which contains approximately 15×10^6 neurons and 10^{11} synapses. These numbers are nearly constant throughout the brain, no matter what the thickness of the cortex (1–3 mm is typical). Primate visual cortex is an exception in that it has over twice the density of neurons as other cortex.

1.2.2 Software

Unfortunately, merely having a million-processor neurocomputer will not be sufficient to implement the new information processing technology. It is also necessary to understand how information may be represented and processed flexibly. The representation and processing of information is essentially a problem of *software*. Judging from the number of applications of neural networks now in progress, one might assume that we have a basic understanding of neural information processing, but this is not the case. A few examples will illustrate the extent of our ignorance.

First, we do not understand how the brain learns. The learning algorithm most commonly used in neural net applications (back-propagation) is almost surely not the algorithm the brain uses, and that is good, because back-propagation is very slow. The slight information we do have about learning in the brain indicates that it uses a rule (Hebb's Rule) that is generally considered too weak for most practical applications. So it seems the brain is using an algorithm that neural network researchers have abandoned as nearly useless. If this is the case, then we need to find out how the brain is able to do this, so that our neurocomputers can learn more quickly.

Second, the are many unsolved problems in the representation and processing of subsymbolic information, specifically in perception and motor control. Even in such a well-studied system as early vision, there is continuing debate about whether the representational primitives are oriented edges, differences of Gaussians, wavelets, 2D Gabor functions, 3D Gabor functions, or something else (MacLennan, 1991).

Third, we know very little about how the brain processes symbolic information, such as language and logic. As discussed above, the flexibility of human symbolic processing derives from the flexibility of the underlying subsymbolic processes. Yet we know little about how the symbolic computation can be built upon subsymbolic processes, which will limit our ability to implement symbolic processing and the higher cognitive processes in the new information processing technology. Thus we will not have true flexible information processing. (See Section 2 for more on this.)

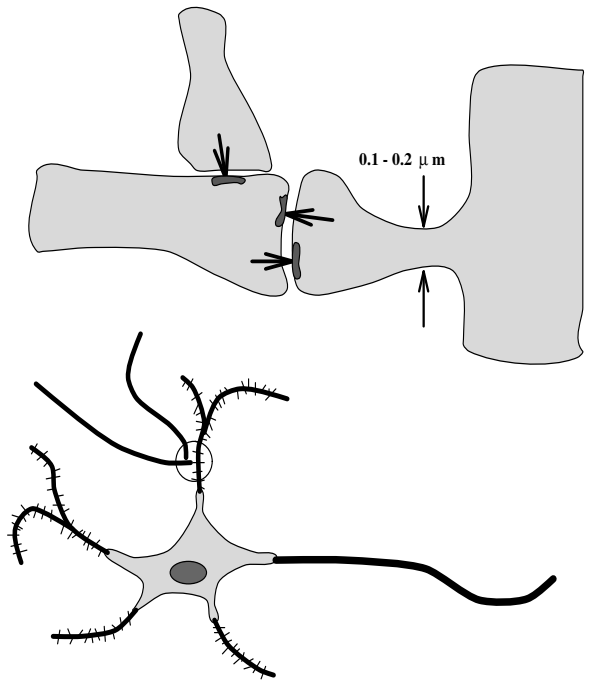

Figure 2: Local Circuits in the Dendritic Net. The lower diagram shows a typical neuron in schematic form. The "hairy" projections on the dendrites represent *dendritic spines*. Many neurons have thousands of dendritic spines, which may be the most basic computational elements of the nervous system. The circled area is enlarged in the upper diagram. An axon coming in from the left makes an excitatory synapse on the dendritic spine on the right. The dendritic spine in turn make an inhibitory synapse *back* onto the axon terminal, thus setting up a negative feedback loop. The synapse from the axon entering at the top modulates the action of the first axon.

Finally, even our understanding of neurons is not as complete as one would suppose from the neural network literature. For example, the model of neurons that is almost universally used treats the dendrites like simple input wires. There is evidence, however, that the dendritic net may be the site of most of the brain's information processing (Shepherd, 1978; MacLennan, submitted). (See Fig. 2.) Furthermore, most neural net models treat neurons as though they change state discretely, whereas it is possible that the interference of continuously varying electrochemical waves is essential to the brain's pattern recognition capabilities (Pribram, 1991). Crick and Asanuma (1986) have noted other important differences between neurons in the brain and in our models. (See Fig. 3.)

The preceding examples show that there are large gaps in our understanding of the brain, which may limit our ability to develop the new information processing technology. It is certainly premature to assume that we know all that is necessary to implement flexible information processing. We conclude that continuing basic research in cognitive psychology and neuroscience should

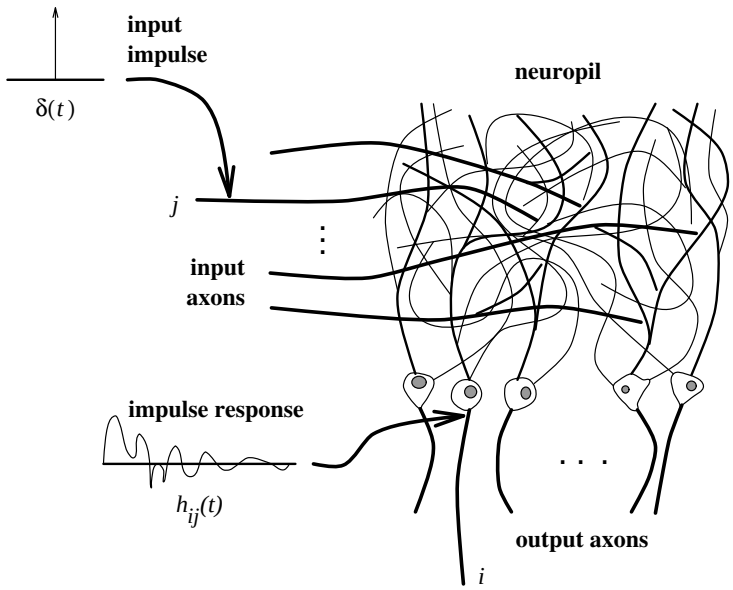

Figure 3: Alternate view of neural information processing based on Pribram's neural wave equation. The neuropil is viewed as a very large number of electrically and chemically coupled oscillators. Incoming impulses drive the dendritic net, causing various resonances. Currents are summed in the dendritic trunks, which act like probes embedded in the neuropil. Currents reaching the cell body and axon hillock may trigger impulses that are transmitted to other areas.

be an essential part of the NIPT project.

1.2.3 Theory

The new information processing technology will require a sufficiently theoretical understanding of the brain so that we will know how to abstract away from the details of neurobiology and be able to implement brain-style flexible information processing in other media. Such a theory should provide a theoretical framework in which to understand the emergence of high level structures from low level interactions, and the control of these interactions by those structures. This theory will constrain and guide research in both the hardware and software areas. In hardware it will show us the kinds of low level functions that need to be implemented to permit higher order structures. In software it will help us understand the information representations that will permit flexible, efficient processing.

1.3 Emergent Computation Research Laboratory

Research into the foundations of flexible information processing is the principal goal of the *Emergent Computation Research Laboratory* currently under development at the University of Tennessee, Knoxville. This research is divided into three areas: (1) theory of emergent computation, (2) information representation and processing (i.e., software), and (3) implementation (i.e., hardware).

The term *emergent computation* has been used for the study of how high-level representation and processing of information can emerge from the interactions of large numbers of simple processes (Forrest, 1990). We believe that a deeper understanding of emergent computation will be necessary for the full exploitation of the new information processing technology, and so we have been investigating several different aspects of emergent computation, including field computation, continuous simulated annealing, continuous spatial automata, and the evolution of communication in populations of simple machines (MacLennan, 1987, 1989, 1990b, 1992; MacLennan & Burghardt, submitted). We are continuing these investigations, and intend to pursue other models from nature that will increase our understanding of emergent computation.

We have been investigating information representation and processing at both the subsymbolic and symbolic levels, and have concluded that there is a need for a fundamentally new framework for understanding the representation and processing of symbolic information. Although we have some preliminary results (MacLennan, in press-a, in press-b), much work remains to be done. We are also investigating the representation and processing of subsymbolic information, including perceptual and motor images (MacLennan, 1991, submitted).

Effective neurocomputation presumes a very large number of processors; we have set a million to a billion as our goal. We do not expect traditional analog VLSI to be able to provide the necessary densities for quite some time, so we have concentrated on optical, opto-electronic, and opto-molecular processes.

These are made more attractive by the fact that field computers (MacLennan, 1987, 1989, 1990a) do not require regular arrays of processors, which should simplify fabrication.

1.4 Conclusions

Achieving the goals of the "information network society of the 21st century" will indeed require a "new information processing technology." However, developing such a technology will depend on continuing basic research in a number of areas, including cognitive science, neurophysiology, epistemology and the theory of emergent computation. Further, research should be directed to the implementation of a neurocomputer with a million to a billion slow-speed, low-precision analog processors. We are optimistic that work in this area will lead to a revolution in flexible information processing.

1.5 References

Changeux, Jean-Pierre (1985). *Neuronal Man: The Biology of Mind.* New York, NY: Oxford Univerity Press.

Crick, F., and Asanuma, C. (1986). Certain Aspects of the Anatomy and Physiology of the Cerebral Cortex. In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2: Psychological and Biological Models*, James L. McClelland, David E. Rumelhart, and the PDP Research Group, Cambridge, MA: MIT Press, 1986, 333–371.

DARPA (1988). *DARPA Neural Network Study*, Fairfax, VA: AFCEA.

Feldman, J. A., and Ballard, D. H. (1982). Connectionist Models and their Properties. *Cognitive Science*, **6**, 205–254.

Forrest, S. (1990). Emergent Computation: An Introduction. In: *Emergent Computation*, edited by Stephanie Forrest, Amsterdam: North-Holland, 1–11.

Ishii, T. (1991). Toward New Information Processing Technologies. *International Symposium on New Information Processing Technologies '91*, Tokyo, March 13–14, 1991, 3–6.

MacLennan, Bruce J. (1987). Technology Independent Design of Neurocomputers: The Universal Field Computer. *Proceedings, IEEE First International Conference on Neural Networks*, San Diego, June 21–24, 1987, Vol III, 39–49.

MacLennan, Bruce J. (1988). Logic for the New AI. In: *Aspects of Artificial Intelligence*, J.H. Fetzer (ed.), Kluwer Academic Publishers, 163–192.

MacLennan, Bruce J. (1989). Continuous Computation: Taking Massive Parallelism Seriously. Poster presentation, Los Alamos National Laboratory Center for Nonlinear Studies 9th Annual International Conference, *Emergent Computation*, Los Alamos, May 22–26, 1989. See also University of

Tennessee, Knoxville, Computer Science Department technical report CS-89-83, June 1989.

MacLennan, Bruce J. (1990a). Field Computation: A Theoretical Framework for Massively Parallel Analog Computation, Parts I–IV. University of Tennessee, Knoxville, Computer Science Department technical report CS-90-100, February 1990.

MacLennan, Bruce J. (1990b). Continuous Spatial Automata. University of Tennessee, Knoxville, Computer Science Department technical report CS-90-121, November 1990.

MacLennan, Bruce J. (1991). Gabor Representations of Spatiotemporal Visual Images. University of Tennessee, Knoxville, Computer Science Department technical report CS-91-144, September 1991.

MacLennan, Bruce J. (1992). Synthetic Ethology: An Approach to the Study of Communication. In: *Artificial Life II: The Second Workshop on the Synthesis and Simulation of Living Systems*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Vol. X, edited by Doyne Farmer, Christopher Langton, Steen Rasmussen, and Chuck Taylor. Redwood City: Addison-Wesley, pp. 631–658.

MacLennan, Bruce J. (in press-a). Continuous Symbol Systems: The Logic of Connectionism. In: *Neural Networks for Knowledge Representation and Inference*, edited by Daniel S. Levine and Manuel Aparicio IV, Hillsdale, NJ: Lawrence Erlbaum, in press.

MacLennan, Bruce J. (in press-b). Characteristics of Connectionist Knowledge Representation. *Information Sciences*, to appear.

MacLennan, Bruce J. (submitted). Information Processing in the Dendritic Net. (Invited presentation, *Second Behavioral and Computational Neuroscience Workshop* (May 18–20, 1992), Georgetown University, Washington, DC.)

MacLennan, Bruce J., & Burghardt, Gordon M. (submitted). Synthetic Ethology and the Evolution of Cooperative Communication.

Pribram, Karl H. (1991). *Brain and Perception: Holonomy and Structure in Figural Processing.* Hillsdale: Lawrence Erlbaum.

Shepherd, Gordon M. (1978). Microcircuits in the Nervous System. *Scientific American*, **238**, 92–103.

Shepherd, G. M. (Ed.). (1990). *The Synaptic organization of the brain*, third edition. New York NY: Oxford University Press.

Smolensky, Paul (1988). On the Proper Treatment of Connectionism. *Behavioral and Brain Sciences*, **11**, 1–74.

2 The Emergence of Symbolic Processes From the Subsymbolic Substrate[4]

2.1 Human Symbolic Cognition

Human symbolic cognition, by which we mean the use of language, logic and explicit reasoning, is much more flexible than symbolic AI systems (MacLennan, 1988). If an expert system has too few rules, then it exhibits *brittle* behavior, failing in catastrophic ways when faced with novel situations or minor exceptions to the rules. On the other hand, trying to anticipate all the situations and exceptions that may occur leads to a proliferation of rules and an exponential explosion in the machine resources required. Human cognition does not have these limitations. When people reason explicitly, their use of categories is sensitive to the context of the problem, and their inferential processes are generated from and constrained by relevance to the situation. This seems to be a natural result of the implementation of these processes, rather than a result of special context rules or relevance rules. Human language use has similar characteristics: it is flexible, context-sensitive and controlled by relevance. Further, flexibility and "softness" is the natural state of both reason and language; "hard" logic and precise language use are special skills that are not easily acquired, and they are special tools used by the expert when they are called for, but not otherwise. How can we achieve similar flexibility in our computer systems?

It is not sufficient to replace categories with hard boundaries by categories with fuzzy boundaries (as is done in fuzzy set theory). Although fuzzy categories do eliminate some sources of brittleness, they do not address the complex processes by which categories may be sensitive to the global context. To achieve the flexibility of human cognition and true context-sensitive symbol use, a more radical reinterpretation of symbolic processing is necessary. We must see context-sensitive holistic processes as the normal mode of operation, and context-free discrete symbol manipulation as a specialized modification of this norm.

These observations have practical implications, for they mean human-like flexibility and competence is unlikely to result from a simple hybrid of neural network technology and expert systems technology. Wherever it appears in a hybrid system, discrete, formal, context-free symbol manipulation will be a source of brittleness and other forms of unskillful behavior. If we want implementations of symbolic cognition that exhibit the flexibility and competence of people's, then they must be built on a foundation that is fundamentally continuous, holistic and context-sensitive.

[4]Text of invited panel presentation, *International Symposium on New Information Processing Technologies '91*, Tokyo, March 13–14, 1991.

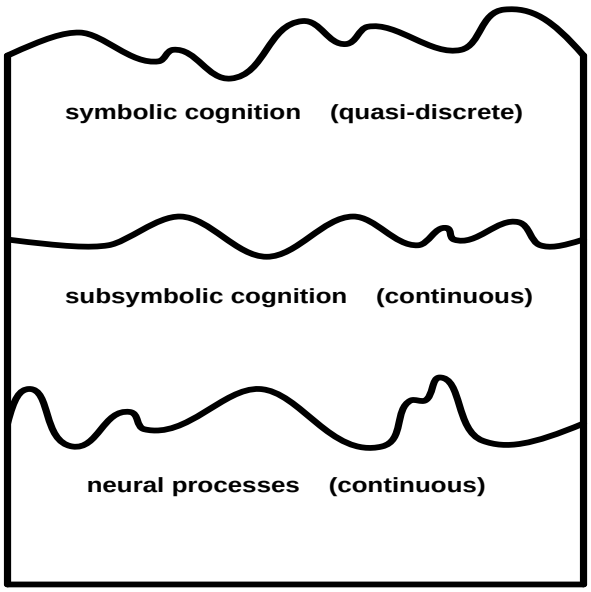

Figure 4: Emergence of Symbolic Cognition from Subsymbolic Substrate. The underlying neural processes of the brain are continuous, as are subsymbolic cognitive processes, including perception, association, sensory-motor coordination, and judgement of similarity. Approximate discreteness appears only at the symbolic level, which may exploit the underlying continuous processes for the sake of flexibility.

2.2 Emergence of the Discrete from the Continuous

Human symbolic cognition is built upon a *subsymbolic substrate*, which is continuous in its principles of operation, and is the ultimate source of the flexibility of human symbol use (Fig. 4; see also Smolensky, 1988). To see this, observe that the basic neural processes are best described as continuous, since information is represented by continuous quantities such as spiking frequency and membrane potential. Most neurons seem to operate as low precision analog devices. Furthermore, the large number of neurons found in the brains of the higher animals means that functional areas can often be viewed as *spatially* continuous. That is, we can view such an area as a continuum of neurons, rather than as a large number of discrete neurons (MacLennan, 1987). Thus the basic neural processes are both temporally and spatially continuous. This is important from a theoretical perspective, because it means that powerful mathematical tools may be brought to bear on the problem of mental representation (e.g., see MacLennan, 1991 and submitted).

In addition, many of the most basic cognitive processes, such as perception, association, sensory-motor coordination, and judgement of similarity, are by their nature continuous. We share these faculties with the lower animals, and we observe that they show the same flexibility in their use as we do. Discreteness is most apparent when we come to higher cognitive processes, such as language use and explicit reason, but they partake of the flexibility of the underlying

continuous processes.

These considerations suggest that a critical research goal is to understand the processes by which discrete and approximately discrete symbols can emerge from continuous processes. Of course this is an inversion of the usual situation in computer science and even logic, where discrete symbols are assumed to be primitive, and continuous quantities are approximated by discrete structures. We must understand how the discrete structures found in logic and language can emerge from continuous representations and processes (MacLennan in press-b, in press-a).

A possible basis for the required theory can be found in the mathematics of Hilbert spaces, which establish a fundamental link between discrete and continuous phenomena (MacLennan, 1990b; Pribram, 1991). The simplest example is the Fourier series for a square wave; here we have a discrete object being approximated by a superposition of continuous objects. Of course Fourier analysis has been extensively applied to the analysis of the visual and auditory systems (MacLennan, 1991), but it is perhaps surprising that it may inform our understanding of symbolic cognition.

A few more examples may be in order. Since Hilbert spaces also permit multiple representations in linear superposition, we can represent in the same functional area several interpretations of a retinal image, or several parses of a sentence. Further, these superpositions can be processed in parallel by a single neural network (MacLennan, 1990a, p. 674n); see Fig. 5. Linear superpositions can also be used to simultaneously represent several objects, thus avoiding problems with illusory conjunctions of features. Such superpositions suggest simple means for shifting objects in and out of the focus of attention, while retaining the contextual sensitivity of the foreground to the background (Fig. 6).

We are not suggesting that Hilbert spaces are the answer to every question about symbolic cognition; there are many other tools available. For example the topology of connected metric spaces seems a useful basis for a qualitative theory of symbolic cognition (MacLennan in press-b, in press-a). In any case, much more research needs to be done.

2.3 Implications for Sixth Generation Technology

Neural networks and related technologies show great promise as efficient, flexible implementations of lower-order cognitive processes, such as perception, motor control, recognition, association and judgement of similarity. Many applications, however, will depend on implementations of higher-order cognitive processes. Unfortunately, efficient, flexible implementations of higher-order cognitive faculties, such as natural language and explicit reasoning, will require an understanding of their relation to the lower-order processes. Although many connectionist researchers are working on this problem, it is still unsolved. Nevertheless, a theory of the emergence of symbolic processes from the subsymbolic substrate will be of great practical and theoretical importance. We conclude

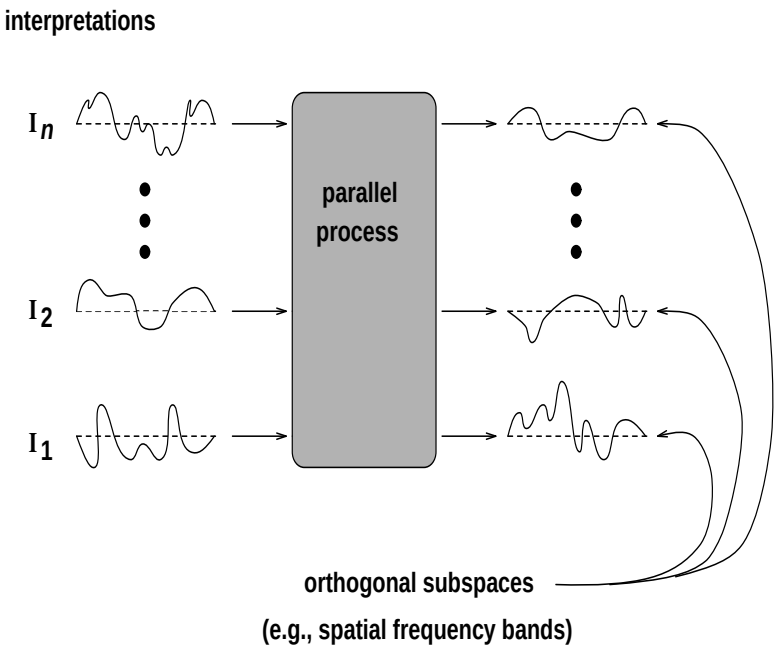

Figure 5: Parallel Processing by Linear Superposition in a Neural Network. Multiple interpretations $I_1, \ldots, I_n$ can be represented in the same neural tissues if they are restricted to (spatially) orthogonal subspaces, i.e., spatial frequency bands. Under these conditions the interpretations I_k can be processed independently and in parallel by a single neural network without leading to illusory conjunctions of features.

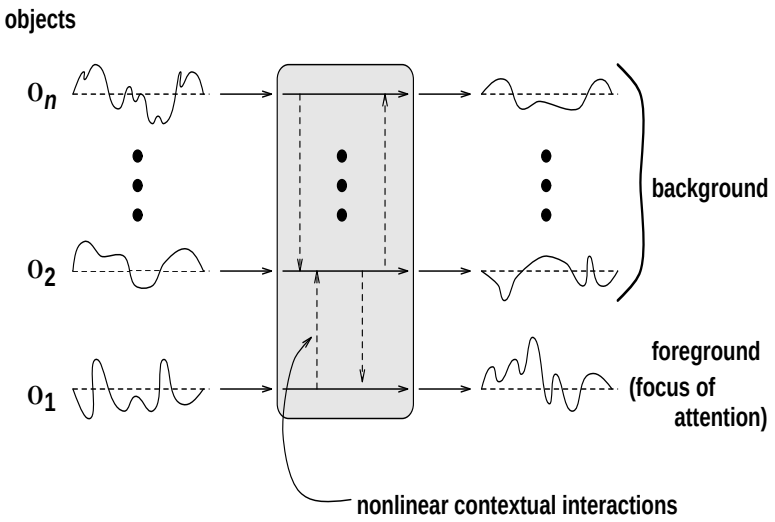

Figure 6: Contextual Effects Represented by Nonlinear Interactions. In some circumstances we may not want independent processing in linear superposition. For example, if the several objects $O_1, \ldots, O_n$ in a scene are represented in linear superposition, then we expect nonlinear interactions between the foreground or focal object O_1 and the background objects $O_2, \ldots, O_n$, since the foreground and the background mutually define each other. Such interactions could, for example, cause a shift of objects between the foreground and background.

with a quotation:

> The disadvantage of regarding things in separate parts is that when one begins to cut up and analyze, each one tries to be exhaustive. The disadvantage of trying to be exhaustive is that it is consciously (mechanically) exhaustive. One goes on deeper and deeper, forgetting to return and sees a ghost. Or one goes on and imagines he has got it, and what he has got is only a carcass. For a thing which retains its substance but has lost the magic touch of life is but a ghost. *Only one who can imagine the formless in the formed can arrive at the truth.*
>
> — Soshi

2.4 References

MacLennan, B. J. (1987). Technology-independent design of neurocomputers: The universal field computer. In M. Caudill & C. Butler (Eds.), *Proceedings, IEEE First International Conference on Neural Networks* (Vol. 3, pp. 39–49).

MacLennan, B. J. (1988). Logic for the new AI. In J. H. Fetzer (Ed.), *Aspects of artificial intelligence* (pp. 163–192). Dordrecht, Holland: Kluwer Academic Publishers.

MacLennan, B. J. (1990a). The discomforts of dualism. *Behavioral and Brain Sciences*, **13**, 673–674.

MacLennan, B. J. (1990b). *Field computation: A theoretical framework for massively parallel analog computation, parts I–IV* (Report No. CS-90-100). Knoxville, TN: University of Tennessee, Department of Computer Science.

MacLennan, B. J. (1991). *Gabor representations of spatiotemporal visual images* (Report No. CS-91-144). Knoxville, TN: University of Tennessee, Computer Science Department. Also submitted for publication.

MacLennan, B. (in press-a). Continuous symbol systems: The logic of connectionism. In Daniel S. Levine and Manuel Aparicio IV (Eds.), *Neural networks for knowledge representation and inference*. Hillsdale, NJ: Lawrence Erlbaum.

MacLennan, B. J. (in press-b). Characteristics of connectionist knowledge representation. *Information Sciences*, to appear.

MacLennan, B. J. (submitted). Information processing in the dendritic net. (Invited presentation, *Second Behavioral and Computational Neuroscience Workshop* (May 18–20, 1992), Georgetown University, Washington, DC.)

Pribram, K. H. (1991). *Brain and perception: Holonomy and structure in figural processing.* Hillsdale: Lawrence Erlbaum.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, **11**, 1–74.