Flexible Computing in the 21st Century*

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

March 13, 1995

1 The New Information Processing Technology

1.1 Introduction

In Tokyo on March 13–14, 1991, MITI sponsored the *International Symposium on New Information Processing Technologies '91*, or *NIPT '91.* 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.

1.2 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

*Revision of article appearing in *Vector Register*, 1991, Institute for Supercomputing Research, 15F INUI Bldg. Kachidoki, 1-13-1 Kachidoki, Chuo-ku, Tokyo 104, Japan.

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, 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.

1.3 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": 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, including 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.

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.

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 has about a million neurons in its brain, and the human brain is estimated to comprise at least a billion neurons. 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.

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, 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.

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). 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.

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 be an essential part of the NIPT project.

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.

3 Emergent Computation Research Laboratory

Research into the foundations of flexible information processing is the principal goal of the *Emergent Computation Research Laboratory* currently being set up 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). 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 1991a, 1991c), much work remains to be done. We are also investigating the representation and processing of subsymbolic information, including perceptual and motor images (MacLennan 1991b).

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.

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.

5 References

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.

Forrest, S. (1990). Emergent Computation: An Introduction. In: *Emergent Computation*, edited by Stephanie Forrest, Amsterdam: North-Holland, 1990, 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. (1991a). 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. (1991b). Gabor Representations of Spatiotemporal Visual Images. University of Tennessee, Knoxville, Computer Science Department technical report CS-91-144, September 1991.

MacLennan, Bruce J. (1991c). Characteristics of Connectionist Knowledge Representation. University of Tennessee, Knoxville, Computer Science Department technical report CS-91-147, November 1991. Submitted for publication.

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, 1992, pp. 631–658.

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

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

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