5p.

SEMIANNUAL STATUS REPORT
October 1985 - March 1986

IN-30236

Contract No. NAS 2-11530
RIACS Task No. 30

THE APPLICATION OF A SPARSE, DISTRIBUTED MEMORY TO THE DETECTION, IDENTIFICATION, AND MANIPULATION OF PHYSICAL OBJECTS

Pentti Kanerva

Research Institute for Advanced Computer Science
Mail Stop 230-5
NASA Ames Research Center
Moffett Field, CA 94035

ACTIVITIES

This was the first half-year of the sparse, distributed memory project. During that time, Kanerva worked on the task full time, and Raugh (RIACS core) contributed significantly by making numerous contacts outside RIACS and by being available to discuss all aspects of the work. Progress was made in four areas: (1) background studies, (2) research planning, (3) professional contacts, and (4) other research-related activities.

1. BACKGROUND STUDIES

To determine the relation of the sparse, distributed memory to other architectures, a broad review of the literature was made. A variety of names is used for these architectures: associative memories, parallel distributed processing (PDP) architectures, connectionist models, and artificial neural nets. We call them PATTERN MEMORIES because they work with large patterns of features (high-dimensional vectors), and we call computers based on them PATTERN COMPUTERS. A pattern is stored in a pattern memory by distributing it over a large number of storage elements and by superimposing it over other stored patterns, and a pattern is retrieved by mathematical or statistical reconstruction from the distributed elements. An important property of pattern memories is that their addressing need not be exact: An approximate retrieval cue will initiate the retrieval of a stored pattern.

N87-12169

(NASA-TM-89231) THE APPLICATION OF A SPARSE, DISTRIBUTED MEMORY TO THE DETECTION, IDENTIFICATION AND MANIPULATION OF PHYSICAL OBJECTS Semiannual Status Report, Oct. 1985 - Mar. 1986 (NASA) 5 p CSCL 09B G3/60

Unclas
44673

For purposes of comparison, three groups of pattern memories will be distinguished according to the topology of their storage elements (adjustable weights). A complete characterization of these memories would include, in addition, a description of their decision elements (threshold units) and of the rules for updating the memories.

GROUP 1. Of the many architectures, the sparse, distributed memory is most like the cerebellar model of Marr and the Cerebellar Model Arithmetic Computer (CMAC) of Albus. Common to them is structural similarity to the random-access memory (RAM) of a computer, typified by arbitrarily many addressable storage locations. These memories are characterized by

$$m \times n'$$

adjustable weights, where m is the number of (n' element) storage locations and n' the number of components in the output patterns (output dimension). These memories are addressed by n-component input patterns (input dimension = n), and any given read or write operation involves but a small subset of the storage locations.

GROUP 2. The associative-memory models of Anderson, Hopfield, Kohonen, Willshaw, and others form a second group. The models are similar to one another in one significant respect, setting them apart from the first group: The number of (n' element) storage locations equals the number of components in the input patterns. Therefore, the capacity of these memories is tied to the size of the input (and output) patterns, whereas in the first group it is not. Furthermore, any read or write operation involves much or all of the memory. These memories are characterized by

$$n \times n'$$

adjustable weights, where n and n' are the input and output dimensions, as above, and usually $n = n'$.

GROUP 3. The multilayer or hidden-unit models of Hinton, Huberman, Kohonen, McClelland, Rumelhart, Sejnowski, and others form a third group. A typical memory in this group is a cascade of two or more memories of the second group. If the "hidden" intermediate layers have $p, q, r, \ldots, z$ components, in that order, the memory is characterized by

$$n \times p + p \times q + q \times r + \ldots + z \times n'$$

adjustable weights. The intermediate layers make it possible to vary memory capacity independently of the sizes of the input and output patterns (of n and n'). As with the memories of the second group, a read or write operation tends to involve much (or all) of memory. More complicated multilayer memories would have more complicated interconnections, but their properties are not generally understood.

GROUPS 1 AND 3. A memory of the first group can, in fact, be realized by a two-layer memory of the third group, with the middle layer consisting of m components. According to the description of Group 3, such a memory has $n \times m + m \times n'$ adjustable weights. When this memory is simplified by fixing the first $n \times m$ weights, we get a sparse, distributed memory of the first group with $m \times n'$ adjustable weights.

GROUPS 1, 2, AND 3. To summarize, memory capacity is independent of the input and output dimensions in the first and third groups and highly dependent in the second (single-layer) group. The memories in the third group are more general than in the first, and therefore possibly more powerful. However, the greater simplicity of the first group pays off in speed of learning: Experiments indicate that it takes at least 100 times as many trials to train memories of the third group as it does to train those of the first. The sparse, distributed memory belongs in the first group and therefore it learns rapidly and can be made arbitrarily large without changing its input and output dimensions. The memory requirements of several tasks were estimated and are summarized in Table 1.

TABLE 1

Realizing Sparse, Distributed Memory in Different Kinds of Hardware (Estimates)

Hardware	Dimension, n	Number of locations, m	Cycles per second	Task
Dedicated DEC 2060	128	10,000	.2-1	Demonstrate convergence properties of the memory
32-node Intel iPSC	128	50,000	1-5	Simple learning by trial and error
16K-processor Connection Machine	200	60,000	50-200	Word parsing in compacted text
Prototype	256	100,000	10-100	Word parsing in compacted text and possibly in speech
Present VLSI potential	1,000	100,000,000	1,000	Language understanding (?)

2. RESEARCH PLANNING

Major effort went into preparing a research plan for the extensive study of sparse, distributed memories and of systems based on them. The plan calls for a three-year study of the architecture in collaboration with Professor Michael Flynn at Stanford and Professor Terrence Smith at UC Santa Barbara. According to this plan, the engineering design of the memory would be studied at Stanford and a resulting hardware-prototype memory would be placed at RIACS, the uses of the memory in controlling a TV camera and a robot arm and a hand would be studied at UC Santa Barbara, and the mathematical properties of the memory and applications at large would be studied at RIACS. Raugh and Kanerva visited funding agencies in Washington, D.C., (NASA, DARPA, ONR) to discuss the proposed research.

3. PROFESSIONAL CONTACTS

Special effort was made to make the project known in the research community and to establish working relations with the community. This was done by participating in C-SAR activities and giving presentations as listed later. In addition, important visits were made with the following people:

Dr. James Albus, Chief of the Robot Systems Division at the National Bureau of Standards. Albus' Cerebellar Model Arithmetic Computer (CMAC) is in the first group of models described above, and his BRAINS, BEHAVIOR, AND ROBOTICS (Peterborough, N.H.: BYTE Books of McGraw-Hill, 1981) charts a new course for the study of artificial intelligence.

Dr. Robert Hecht-Nielsen, director of the DARPA ADAPT program, and Drs. Jack Smith and Allen Wu at the Rancho Carmel AI Center of TRW, San Diego. They are developing the Mark IV computer for modeling artificial neural nets and are now preparing to build a sparse, distributed memory as well.

Professor David Rumelhart and Dr. Ronald Williams at the Institute for Cognitive Science at UC San Diego. They (Rumelhart and McClelland) have edited a major two-volume book on Parallel, Distributed Processing (PDP) architectures, being published in 1986.

Professor Terrence Smith and his graduate students Gilbert Pitney and Umesh Toylekar at the Computer Science Department of UC Santa Barbara. They have programmed a simulator of the sparse, distributed memory on an LMI Lisp machine, for the study of robot manipulatory control with visual feedback.

The most fruitful cooperation has been with Professor Michael Flynn and his graduate student, Bahram Ahanin, at the Electrical Engineering Department at Stanford. It is leading into the development of a prototype memory for the study of design issues and for application studies. A preliminary hardware design was made by Mr. Ahanin and is referred to in Table 1 as the 'Prototype'.

4. OTHER RESEARCH-RELATED ACTIVITIES

Kanerva attended a one-week workshop on the Intel iPSC Hypercube computer at the Intel Scientific Computing Center in Beverton, Oregon, in October 1985. RIACS has a 32-node iPSC computer, which will be programmed to simulate the sparse, distributed memory.

Kanerva attended two meetings of the Consortium of Space Automation and Robotics (C-SAR) at UC San Diego, one with Raugh and the other alone. Both organizational and research issues were discussed in these meetings. The sparse, distributed memory project could contribute to C-SAR activities by helping to identify promising research topics and by suggesting approaches to some of the selected topics. The memory model can shed light on two issues in particular: How does an autonomous system build an internal model of the world, and how might telerobots be used as a way to develop more and more autonomous robots?

PUBLICATIONS

Kanerva, P. PARALLEL STRUCTURES IN HUMAN AND COMPUTER MEMORY (Rep. No. TR-86.2). Moffett Field, CA: RIACS at NASA Ames, January 1986.

ABSTRACT. If we think of our experiences as being recorded continuously on film, then human memory can be compared to a film library that is indexed by the contents of the film strips stored in it. Moreover, approximate retrieval cues suffice to retrieve information stored in this library: We recognize a familiar person in a fuzzy photograph or a familiar tune played on a strange instrument. This paper is about how to construct a computer memory that would allow a computer to recognize patterns and to recall sequences the way humans do. Such a memory is remarkably similar in structure to a conventional computer memory and also to the neural circuits in the cortex of the cerebellum of the human brain. The paper concludes that the frame problem of artificial intelligence could be solved by the use of such a memory if we were able to encode information about the world properly.

PRESENTATIONS

Raugh, M. "Kanerva's Sparse, Distributed Memory: A RIACS Project." NASA Ames SETI Project, November 1985; Xerox PARC, January 1986.

Raugh, M. "An Introduction to RIACS and Kanerva's Sparse, Distributed Memory." Bolt, Beranec, and Newman, Cambridge, Mass., November 1985; Thinking Machines, Cambridge, Mass., November 1985; KLA, Santa Clara, January 1986.

Kanerva, P. "The Organization of an Adaptive System Based on Sparse, Distributed Memory." UC Santa Barbara Computer Science Colloquium, February 1986.

Kanerva, P. "BRAINS, BEHAVIOR, AND ROBOTICS by James S. Albus." Stanford CSLI TINLunch, March 1986.