

PREFACE

This volume brings together recent research contributions of ten authors to the field that is variously called neural networks, connectionism, parallel distributed processing, or biological information processing. All of the articles were carried out through collaborative work with members of the Center for Adaptive Systems (CAS) at Boston University.

The CAS carries out two types of theoretical activities which, although conceptually independent, have turned out to be mutually reinforcing in our case. One type of activity studies the fundamental design principles and mechanisms needed to explain and predict large data bases about brain and behavior. The other type of activity generates novel architectures for implementation as intelligent machines in technological applications.

There is no necessary reason why all studies of biological intelligence should go technological, or conversely. To validly advance understanding of biological intelligence, one must explain and predict lots of biological data. To advance understanding of machine intelligence, one must solve outstanding technological problems. Diminishing returns set in when a flimsy technological advance is propped up by saying it works just like the brain, or when a metaphorical brain theory devoid of data implications is heralded as the next hi-tech sensation.

Why, then does the type of research done at CAS lend itself to both biological and technological applications? The research at CAS is grounded in its sustained analyses of biological data. This research has proved to be relevant to technology both because of the types of problems we study and the methods that we use to solve them.

In particular, we study problems requiring real-time adaptive responses of individuals to unexpected changes in complex environments. These are the types of problems that humans and mammals need to solve in order to survive. These are also among the types of technological problems that traditional scientific and engineering approaches have not already well-handled.

Our methods for attacking such problems are systematic and rigorous. We typically begin by analysing a huge interdisciplinary data base within a prescribed problem area. In our work on preattentive vision, for example (Chapters 1-4), we have studied data from many parts of the vision literature—data about boundary completion, texture segmentation, surface perception, depth perception, motion perception, illusory figures, stabilized images, hyperacuity, brightness and color paradoxes, multiple scale filtering, and neurophysiology and anatomy from retina to prestriate cortex. Only through the sustained analysis of many hundreds or even thousands of such experiments can one accumulate enough data constraints to discard superficial modelling ideas and to discern a small number of fundamental design principles and circuits.

Such concepts do not make themselves known, however, through a purely bottom-up shifting among huge heaps of data. They come into view by thinking about how these data could arise as emergent, or interactive, properties of a real-time process engaged moment-by-moment by the external environment. Being able to think about an immense mass of static data in terms of the real-time processes that generate these data is an art. It is, I believe, the rate-limiting skill in our work. These real-time processes gradually become discernable through the active confrontation of a huge data mass with known theoretical

principles, mechanisms, and computations about real-time neural network processes to test for matches and mismatches. Through this approach, a series of design paradoxes, or trade-offs, come into view which balance many data and computational requirements against one another. Gradually the accumulated impact of these design tradeoffs creates such an intense intellectual pressure within the emerging scheme of ideas that every fact and hypothesis ramifies through it with multiple implications. By this point, a well-defined family of real-time neural network architectures has usually come into view, supported by a new computational theory that is often quite invisible to a merely passive compilation of the data.

We develop and test these architectures using rigorous mathematical techniques and systematic parametric series of computer simulations in order to gain a complete formal understanding of their emergent, or interactive, properties. The combination of working on problems for which both biology and technology need answers, and developing these answers into rigorously characterized computational structures, makes such work equally applicable to quantitative data analysis and to efficient technology transfer.

This rigorous approach has led to real-time neural network architectures that provide explicit examples of intelligent systems which overcome classical bottlenecks in stability, adaptability, scalability, capacity, and speed that have hampered the further development of AI algorithms. Because the demonstrations that guarantee these properties take the form of rigorous mathematical theorems and parametric computational analyses, they provide a firm foundation upon which software and hardware applications may confidently be supported.

What kind of problem can such a neural network architecture solve? Each architecture is being developed to supply a general-purpose solution within a focussed problem domain—what has been called a solution of a *modal problem*. A modal architecture is less general than a general-purpose digital computer but much more general than a typical AI algorithm.

The modal problems analysed in this volume form part of the answer to a central question about real-time adaptive behavior: How does a freely moving human or robot learn to approach an affectively valued object? Despite the apparent simplicity of this question, its complete answer would require an understanding of at least four types of modal problems. Contributions to analysing all these types of modal problems are represented herein, including problems of:

perception (Chapters 1–4), including architectures for biological and machine vision, which clarify how processes of emergent boundary segmentation and featural filling-in can cooperatively utilize scenic data about boundaries, textures, surfaces, multiple spatial scales, and stereopsis to generate a hyperacute preattentive representation of 3-dimensional form in which variable illumination conditions are discounted;

cognitive information processing (Chapters 5–8), including architectures for adaptive pattern recognition, nonstationary hypothesis testing, self-adjusting parallel memory search, updating of working memory, and automatic reallocation of attentional resources;

cognitive-emotional interactions (Chapters 9–11), including architectures for rapidly focussing attention on environmental events and hypotheses which predict behavioral success based upon prior satisfaction of internal constraints, as in the action of rewards, punishments, homeostatic rhythms, or the unexpected nonoccurrence of expected goals,

leading to new insights concerning how such adaptive mechanisms may generate irrational behavior when confronted with environments requiring decision making under risk; and

goal-oriented motor control and robotics (Chapter 12), including architectures which circumvent classical combinatorial explosions to show how invariant properties of flexible arm trajectories can be generated as emergent real-time properties of nonlinear neural interactions, rather than as explicitly pre-planned commands, and how self-calibration of movement command parameters can be learned automatically after partial accidents or other unexpected environmental feedback.

Although each of these projects can, at least in part, be carried out independently, they can also collectively benefit from efficiencies of cooperation, interfacing, and scale when they are organized as part of a coordinated research program aimed at the design of intelligent machines capable of autonomous adaptive real-time operation in unanticipated environmental situations, as they are at an interdisciplinary institute such as CAS.

In order to partially cope with the interdisciplinary demands that are commonplace in analysing such modal problems, each staff member at the CAS typically has training in at least three of the four fields: computer science, mathematics, psychology, and neurobiology (the 3/4 Rule) so that every individual has both good design intuition and technical ability to carry out advanced modelling. It is really quite futile to ask people who individually know the data, or have strong computational technique, but not both, to understand each other, let alone to collaborate effectively with each other. Key scientists in such collaborations need to individually possess both design intuition and appropriate formal technique. Consequently, despite its relatively small size, Center scientists hold advanced degrees in mathematics, computer science, psychology, neurobiology, engineering, and physics.

Research is carried out in small cooperative groups and each scientist works on multiple projects which often cut across fields. The unifying impact of available neural network theories and methods make such diversity possible, since results discovered in one problem domain often bring new insights to seemingly unrelated problem domains; for example, results about reinforcement mechanisms have shed new light on mechanisms of vision, mechanisms of circadian rhythms, and mechanisms of motor control because all these mechanisms use specialized variants of a general design for opponent processing (fear-relief, red-green, awake-asleep, push-pull) that is called a *gated dipole field* (see Chapters 2, 10, and 11). Such cooperative interdisciplinary research teams may become the norm in future theoretical investigations of intelligent architectures due to the sheer complexity of the scientific and technological problems.

If you ask my CAS colleagues how they got trained to satisfy the 3/4 Rule, you should not, however, expect to get a straightforward reply. The answer may vary from "accidentally" to "against all odds" to "I'm a very unusual person." It simply is not easy to get systematic training in the major data domains and computational techniques that are needed to do advanced theoretical work on mind and brain topics. That is why several universities, including our own, have begun to put together new interdisciplinary programs or degrees to make it easier for students to enjoy learning about this field.

I hope that volumes such as the present one, which bring together research that includes a variety of the problems, models, and methods at the forefront of neural network research, will make it easier for individuals at all levels to acquire the interdisciplinary knowledge needed to appreciate and contribute to this exciting field of scientific inquiry.

The work reported within this volume has been supported by the Air Force Office of Scientific Research, the Army Research Office, the National Science Foundation, and the Office of Naval Research. We are grateful to these agencies for making this work possible.

We are also grateful to Cynthia Suchta for doing a wonderful job of typing and formatting the text.

Stephen Grossberg
Boston, Massachusetts
August, 1987