Preface to the Second Edition

This work seeks to develop and integrate two themes. One is that of associative memory, a topic well worked over, but still needing a development from clear principles. The other is the particular theory of oscillatory networks presented in a highly significant series of papers by W.J. Freeman and his collaborators over more than twenty years; see the references in Chapter 11.

The concern is to investigate artificial neural networks (ANNs) rather than biological neural networks (BNNs) — a stance which absolves a non-biologist of presumption while leaving him full freedom. However, as is usual, one looks to physiology for both inspiration and confirmation; hoping also to repay the debt by the development of some insight from the idealised ANN models.

Part II of the text attempts to develop rational designs for an associative memory from some set of principles. Such principles must stem from an understanding of the function which is required, and this comes first in Chapter 8, on 'fading data'. The function of a full associative memory is to form an inference on the basis of data which may be fleeting and which cannot be stored without degradation. The dynamics of the memory must then be such as to hold significant aspects of the data while the inference is being formed. Quantisation, in some sense, is necessary.

A probability-maximising principle, which turns out to have optimality properties, is adopted to achieve this. It yields an appealing solution: a dynamic, autoassociative and somewhat fuzzy form of the Hamming net. However, the equivalent structure turns out to have real interest, particularly when the images represented by the data are compound in some sense. The possibility of compounding implies a structure which we regard (see Chapter 14) as a fair and insightful model of the transmission/processing loop constituted by the olfactory bulb and the anterior olfactory nucleus in the olfactory system. Of course, one must beware of facile identifications, but the two structures show striking and unforced correspondence on a number of significant features. We find the correspondence striking enough that we venture a couple of testable predictions; see section 14.5.

This investigation also clarifies a number of side issues; e.g. it explains how the 'spurious' equilibria of the Hopfield net have virtually been designed into it.

Readers who wish to follow just this associative-memory theme can probably get away with reading Chapters 5, 6, 8 and 14. Readers determined to plunge *in medias res* should make what they can of Chapter 14.

The other theme is that of oscillatory operation, so evidently a feature of biological neural systems. This must convey some advantage, the obvious one being that, since absolute activity levels are almost meaningless in the biological context, information must be transmitted by a variable rather than a static signal. An oscillatory signal can then be seen as a way of transmitting information despite a background of both noise and irrelevant slow variation in baseline activity. However, one-way 'transmission' is far too naive a concept; biological systems are locked into a dynamic whole by both forward and backward connections, and achieve communication by entering a joint dynamic state.

No approach to these ideas has been more insightful and thorough than that of Freeman and his school, which has combined anatomical observation and physiological experiment with hard thinking and mathematical analysis. In this way Freeman has elucidated a number of fundamental mechanisms and deduced very convincing models (notably for the olfactory system) tested by analogue simulation.

However, while the Freeman theory produces a physiologically faithful model for an oscillator which has certain quite specific properties as a threshold element, and shows how these can be coupled together to form systems of the character and behaviour which are observed, the particular structures suggested for the associative-memory function do not go beyond a Hopfield net. Again, design principles are needed.

In Part III we try to supply these by incorporating the Freeman mechanisms into the nets derived in Part II. Both sets of models are dynamic, but the static equilibria of the Part II models are now replaced by dynamic equilibria. The combination of oscillatory units and the global standardising operation inherent in the Hamming net has a remarkable consequence: the generation of a slow square-wave global oscillation, the 'escapement oscillation', which explains the occurrence of gamma-wave bursts and has a clearly synchronising effect.

Readers who wish to follow just the oscillation theme can read Chapters 11 and 12 and much of Chapter 14 directly, but can probably not read Chapter 13 and complete Chapter 14 without cutting back to Chapters 5 and 8.

Finally, we should clarify the reference to chaos in the title. Any electroencephalogram has a truly chaotic look about it, but we regard systems as operating in spite of chaos rather than in virtue of it. It is a matter of observation that the major components of the olfactory system show relatively simple behaviour in isolation — either a stable equilibrium or a limit cycle — but that when coupled together they show the complex dynamics which one would term chaotic. This is true both of the actual organism and of Freeman's electronic analogue. The appearance of chaos seems to be a consequence of complexity plus transmission delays in the total system. However, the point, made well and early by Freeman, is that it does not matter. If the system is locked into a synchronised whole then communication is achieved by the sensing of the joint dynamic metastate.

I am most grateful to Colin Sparrow, who was not only good enough to carry through the calculations of Chapters 11–13, but also provided steady expertise and perceptiveness. The launch of the first edition in 1998 was marred by a marketing miscalculation; I am most grateful to Imperial College Press for giving the venture a second and better chance. I have made no major changes to the text but, on reading it after this interval of time, am convinced that the reader would do well to follow the short-cuts suggested above. That is, that a reader interested primarily in the associative memory theme could well proceed directly to Chapters 5, 8 and 14; one interested primarily in the oscillation theme could proceed directly to Chapters 11, 12 and 14, but then supplement these by a return to Chapters 5, 8 and 13.

The three graphs of the cover design represent, in descending order: a voltage trace from the olfactory bulb of a rat, a trace from Freeman's electronic simulator of the olfactory system, and a trace from a two-stage Freeman oscillator with feedback and dynamic stabilisation of signal strength.