

Chapter 1

Introduction and Aspirations

Readers who like to skip introductions are advised to read the preface, which also sets out our two central aspirations. One is to develop rational designs for an associative memory from a clear set of principles; principles which must in turn stem from an understanding of required function. The other, seemingly unrelated, is to study the role and effect of oscillatory operation in neural networks, leaning heavily upon the many insights brought to this topic by W.J. Freeman. An understanding of rational design is at least as necessary for this case as for any other; hence the connection.

Neural nets, whether biological or artificial, are concerned with the calculations that go on inside a system. These convert information coming in (from a given set of sensory organs) into actions (realised by a given set of motor organs). In the simplest case we regard this as a one-stage procedure. There is a data input y which is converted into an action output u by a rule $u = \phi(y)$. The function ϕ is realised by the neural network and represents the response of the system. One would like to design the network so that this response is the appropriate one, on some criterion.

This response might be the response of a dog to a scent, of a factory manager to an order book or of an investor to market information.

By restricting attention to the one-stage case we of course exclude the fascinating and much wider issues raised by a changing environment, by the possibility that one's actions might affect that environment, or by the possibility of interaction with similar entities. These are indeed higher-order questions, whose resolution must rest upon that of the simplest case.

The response function ϕ is termed a 'decision function' or an 'action rule' in other contexts, and statistical decision theory gives a basis for the optimisation of ϕ . However, this optimisation requires the specification of a full statistical model and optimisation criterion. Furthermore, it gives no

guide to the physical realisation of ϕ and takes no account of the cost of such realisation. It presumably gives an idea of ideal performance, although the optimal solution may be non-robust. That is, performance of the rule deduced may degenerate disastrously if conditions depart even slightly from those assumed. The only protection against this is to optimise average performance over a sufficiently wide class of models.

The essence of the artificial neural net (ANN) approach is that the response rule is improved by evolution under actual operation. The neural net is modified (usually by perturbation of parameters) through the application of a 'learning rule' which improves performance in the light of experience. This offers the prospect of a procedure which is self-optimising under a wide range of assumptions, and so is adaptable, robust and self-righting. It might also give insight into the emergence of structure, and into the functioning of biological analogues. It constitutes the 'ultimate package', in being to a high degree assumption- and analysis-free. However, some ingenuity must be built into the learning rule if this is to be reasonably economical in time and effort. The mutation/selection rules of Darwinian evolution are plainly effective, but levy a vast cost in time and organic material.

However, to characterise operation in terms of a blind response rule is too shallow, even for the simple one-stage case, and this is implicitly recognised in ANN theory. The data y may be regarded as an imperfect indicator of the actual 'state' x of (the environment ('state' at least in relevant respects)). So, the state for the dog above would be a complete identification of the cause of the scent. If one knew what x was, then decision would be easier to the extent that uncertainty would have been eliminated. Under at least some circumstances, and at least approximately, the deduction of an optimal action can be decomposed into the formation of an 'estimate' \hat{x} of x from y , followed by the determination of the action u appropriate to \hat{x} .

The perception of the concepts of state and its estimation could be regarded as the first glimmer of intelligence on the part of the optimiser, although perhaps not yet on the part of the net itself. A device which simply forms this estimate \hat{x} is an *associative memory*, in that it evokes the value \hat{x} of x which is most consistent with the observations y , in some sense. A device which further holds this 'memory trace' in evoked form over time is a *storage memory*. So, time can be relevant, even for a one-stage system. In fact, dynamic effects can enter naturally in quite another way. In seeking

the x which is ‘most consistent’ with observations y we are looking for a value of x which is extremal in some sense. This extremal problem is best solved, not by a pedestrian search through x values, but by a system of dynamic equations which seeks the extremum much more immediately — with much more sense of the geometry of x -space. Beckerman (1997) gives a graphic example of this: a protein model, governed by high-dimensional nonlinear dynamics, finds its minimal-energy configuration in about 10^{-10} of the time that an exhaustive search would take.

In fact, these two reasons for dynamic operation combine in a curious way, which finds expression first in Chapter 8. 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 — what we shall henceforth refer to as ‘fading data’. The dynamics must then be such as to combine the operations of inference (i.e. of search for the best-matching memory trace) and of retention of the significant aspects of the data during this search. In the end, the only aspect of data retained is identification of the best match.

We shall deliberately restrict ourselves to the estimation problem; if one can deal with it, one can probably deal also with the action problem. Even this reduced task is substantial enough; there is of course an enormous literature on associative memories. However, much of this is concerned with the discussion of relatively *ad hoc* proposals, and attention is split between the studies of performance of the net and of performance of learning rules.

We shall also restrict ourselves purely to the design problem. So, central as the learning aspect is to ANN theory, we shall not consider it (save in the survey of Chapter 10) but shall return to classical statistical optimisation principles applied to an assumed model. However, these principles must be relaxed and given dynamic form if the rules deduced are to find natural realisation in a neural net, and also if they are to be able to deal with fading data.

The search for a design principle implies a concern for performance. This concern raises a number of points, one of them being that of compound memory traces; see Chapter 6. It also leads to an attempt to throw yet more light on that already highly exposed object: the Hopfield net. As we argue in Chapter 6, the net is plainly not ‘right’; for one thing, it can deal only with memory traces of a very specific statistical character. Nevertheless, it exerts a fascination; a fascination which we assert stems from the fact that

the net indeed addresses the fading data problem, even if not very explicitly intended to do so.

The fading data problem is introduced in Chapter 8, and resolved by appeal to a probability-maximising principle. The probability-maximising algorithm (PMA) thus derived is almost trivially simple: in effect, a dynamic, autoassociative and somewhat fuzzy form of the Hamming net treated in Chapter 5. It thus has an immediate network implementation, and we show (Chapter 9) that it has optimality properties. Simplicity does not imply triviality; the Hamming net has structure which develops unexpectedly when some degree of compounding of traces is permitted; even more when oscillatory operation is added.

The frequency of reference to biological models and biological reality increases as we progress through the text. This is partly because these constitute what John Taylor, among others, has referred to as the only existence proof we have (of neural nets which function at the level of performance that one would wish). It is then natural to seek guidance from them. Nature demonstrates that neural networks successful in every relevant sense (scale, performance, reliability, adaptability, economy) really are possible, and can indeed show a sophistication infinitely beyond present ANN aspirations.

At a much more specific level, the mathematics seem to indicate that reliable performance at positive information rates can be reconciled with economy of realisation only if one considers compound stimuli (Chapter 6). We are then led to consider the example of the olfactory system, widely studied because of its relative simplicity. This is a good example, in that the odours activating it can plainly be simply superimposed, so constituting additively-compounded stimuli. The visual system is, of course, studied even more, but is greatly more complicated in that it has to cope with spatial structure, non-additive compounding, movement, and the recognition of images subject, not merely to superficial distortion, but to whole families of radical transformations.

The example of the olfactory system is considered explicitly in Chapter 14, where we argue that the net derived from the PMA provides a very plausible model of the transmission/processing loop constituted by the olfactory bulb and the anterior olfactory nucleus. The literature is of course littered with wishful, facile and deluded claims of this type. However, there are striking and unforced correspondences between the two structures on a number of significant features. Indeed, we are led to make testable predictions on the strength of them.

In Part III we follow up the commonplace of observation: that most biological neural nets (BNNs) are oscillatory in their functioning. The obvious reason for such a mode of operation is that suggested in the preface: that absolute activity levels are almost meaningless in the biological context. Information must then 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 an irrelevant slow variation in baseline activity.

Quite the most impressive and thorough investigation of this matter has been that sustained by W.J. Freeman over a number of years (see the references under his name, plus the other references quoted in Chapter 11). Freeman has particularly studied the olfactory system, at the levels on the one hand of anatomy and experiment, and on the other of modelling, mathematical analysis and analogue simulation. He has developed a model of neural masses which retains essential neuronal dynamics, has derived a physiologically faithful model of a fundamental oscillator which has quite specific properties as a threshold element, and has demonstrated that systems can be built up from these components which show the character and behaviour of living systems.

Part III is devoted to development of Freeman's proposals. We set out his general ideas in Chapter 11 and develop some of his models in Chapter 12. However, as mentioned in the preface, the detailed structures Freeman suggests for the associative-memory function do not go beyond the Hopfield net. In Section 12.8 and Chapter 13 we transfer the PMA of Chapter 8 to the Freeman context; some interesting effects emerge. The consequent net is compared with olfactory anatomy in Chapter 14.

Use of the term 'chaotic' in the title may raise associations or expectations which we do not intend. We use it as Freeman does: not as characterised by Hausdorff dimension or by an infinitely fine-grained sensitivity to initial conditions (meaningless in this context!), but by a pattern of oscillation which is complex to the point of apparent irregularity. This complexity is an indication of robustness rather than sensitivity. When one sub-system passes signals to another (e.g. the olfactory bulb to the pyriform cortex) it is not a question of the second sub-system's 'resonating' to a pure carrier frequency of the first. The two constitute components of a joint system, coupled in both directions, and the oscillations are those characteristic of the joint system. Hence both their complexity and yet their coherence between sub-systems. The oscillation is a 'carrier' principally in that its relative amplitude in different parts of the cortex (say) is

characteristic of the external stimulus applied to the system. More exactly, it is characteristic of the class (of a defined set of classes) to which that stimulus belongs.

Both actual organisms and Freeman's electronic analogue show oscillations in neural potential of a character which one would regard as chaotic. This is to be expected; a consequence of the necessary complexity and nonlinearity of these systems and of the inevitable transmission delays. The point is that the systems can be seen as functioning in spite of chaos, rather than in virtue of it.

Organisation of the text

We have already used the only three acronyms which we shall find habit-forming: ANN (artificial neural net), BNN (biological neural net) and PMA (probability-maximising algorithm). In Chapter 14 a number of standard anatomical acronyms are used, and are explained there. Vectors are column vectors, unless otherwise stated.

Sections, equations, theorems and figures are numbered consecutively throughout a chapter, with a chapter prefix. So, equation (3.18) is equation (18) of Chapter 3. However, the chapter prefix is often omitted when reference is made within a chapter.