Preface

Out yonder there was this huge world, which exists independently of us human beings and which stands before us like a great, eternal riddle, at least partially accessible to our introspection.

Albert Einstein

The Brain Metaphor and Neural Networks

The capabilities of even the most powerful computing machines simply pale into insignificance upon conscious reflection of the awesome power of the computation that the human brain performs. With some 100 billion neurons, averaging 10,000 synapses per neuron, the human brain comprises a phenomenal 1000,000,000,000,000 synapses! Computation in the human brain is *non-linear* (in terms of the functions that neurons and their ensembles compute); it is *non-local* (in terms of the parallel nature in which the brain processes information), and *non-stationary* (in the sense of time varying patterns of neuronal activity). Succinctly, the brain is a massively parallel feedback dynamical system.

Can we *replicate* this kind of complexity in artificial systems in order to exploit the computational metaphor that the human brain provides? To some extent, yes; to a large extent, no. The best we can possibly do is to attempt to understand how the brain imparts abilities like *perceptual interpretation*, *associative recall*, *common sense reasoning*, and *learning* in humans, and then develop powerful models and applications that perform with similar functionality.

The artificial neural network paradigm is constantly motivated and constrained by neuronal analogies, rather than by the attempt to model real neurons. Whereas *neural modeling* is concerned with modeling specific neural circuits of real brains, *neural networks* deal with *neurally inspired* modeling [24], where the primary goal is to abstract principles of computation employed by the brain, and ways in which its working can be modified.

For a number of reasons, this is not an easy task. Major areas of the brain are interconnected to one another in a complex anatomy, and computation is distributed over various areas in parallel. It is therefore difficult to separate one function from another. As Carver Mead wrote over a decade ago, "simple neural systems based on clear, obvious principles may once have existed, but they are buried in the sands of time. Billions of years of evolution have presented us with highly efficient, highly integrated, and impossibly opaque systems [382]."

To complicate matters further, detailed models of simple neural systems such as the one found in the extensively studied marine snail *Aplysia Californica*, don't scale up "linearly". Our reductionist assumption—that if we understand in detail the operation of a neuron we will understand the operation of the brain—no longer holds. The complexity of nonlinear computational systems like the human brain stems from the myriad ways in which ensembles of *self-organized* components interact and work in unison. Neural computation is a macroscopic phenomenon that results from the interaction between ensembles of interacting neurons. Unfortunately, such computation is only microscopically evident in the behaviour of any single neuron.

Genesis

From the genesis of my first ideas on neural networks in 1986, through to the completion of this project, the words of Bertrand Piccard[†] quoted at the beginning of the book ring true. It was Most Revered Professor P. S. Satsangi who permitted my foray into the realm of Neural Networks back in 1990 during the course of my thesis. He gave me an uncommon freedom to explore the unknown, and our many ensuing discussions nurtured my ideas on the subject, eventually to evolve over the years into the present text. As I write this Preface, I am grateful for, and aware of the *invisible* support I received, without which you would not have been reading this book.

In 1991 we offered a course on the subject to postgraduate students in the Department of Physics and Computer Science at D.E.I., which evolved to the present PHM 802 Neural Networks. This book has grown out of my lecture notes collected over the past twelve years that I have taught this course.

Scope

Neural networks is a subject that is an integral component of the ubiquitous soft computing paradigm. An in-depth understanding and appreciation of the field requires some background of the principles of neuroscience, mathematics (linear algebra, analysis and dynamical systems), and computer programming. This is a difficult mixture of subjects to deal with. *Neural Networks: A Classroom Approach* aims to achieve a balanced blend of these three areas to weave an appropriate fabric for the exposition of the diversity

[†]Pilot of Brietling Orbiter 3—the first hot air balloon to go non-stop around the world.

of neural network models. There are many fundamental neural network learning algorithms. This book concentrates on the entire repertoire of foundation algorithms of neural computing in a self-contained fashion.

Neural Networks: A Classroom Approach is intended for a first course on neural networks. Almost every university now offers such a course towards the fourth year of an undergraduate engineering program or the first year of a Masters program in Computer Science or Electrical Engineering. This book is suitable for adoption at both these levels—senior undergraduate and graduate level. For any such *course*, this book will serve as the primary text. It is designed to be completely self-contained, and to provide a complete coverage of curriculum requirements for a typical senior undergraduate or graduate level course on neural networks. At the same time, it will also serve as a supporting text for courses on soft computing, artificial intelligence, and neuron modelling, whose entry into university curricula is on the increase. Given the minimum expected level of mathematical expertise required, the book should be readable by non-engineering majors in subjects such as Physics and Mathematics as well. In addition, the book is suitable for researchers who would want to read about how the standard neural networks differ from more neurobiologically oriented next generation models such as pulsed neural networks. The book also contains a detailed up-to-date discussion and review on the role of neural networks in soft computing.

What Makes This Text Different from Others?

The approach followed in this book is unique in style, sequence, and coverage. It stresses on an intuitive and geometrical understanding of the subject. It is based on the premise that an intuitive understanding of the inner working of an algorithm cannot be achieved without a mental picture created from geometry and heuristics, handworked examples, programming, and a firm mathematical treatment.

Geometry and handworked examples build our intuition. Programming cements it. The level of mathematics is carefully balanced throughout so that *no detail is sacrificed*, while ensuring that the reader doesn't lose sight of the "forest for the trees". For each model, the text first builds up a geometrical picture and throws in just enough mathematics as is required to get the neural network algorithm going. The stress is on **heuristic explanations** to provide an understanding of underlying principles. Gradually, the mathematical treatment becomes increasingly firm without sacrificing intuition at any stage. Also, rather than relegating relevant mathematical support to an Appendix, the maths is integrated into the text which makes for a self-contained treatment. Each foundation neural network model is supported by one or more real world applications. Neuroscience findings pepper the text to keep bringing back the reader back to reality. The style of the manuscript is easy and lucid and its title aptly reflects its approach—to put things across as we do with chalk and board.

The approach of this text is distinct from others in that it:

- Provides an understanding of the **underlying geometry** of foundation neural network models while stressing on **heuristic explanations** of theoretical results.
- Works through each neural network algorithm using **handworked example**s.
- Integrates **pseudo-code** operational summaries and well documented **MATLAB code** segments for all models.
- Presents detailed **computer simulations** for all models presented with explanations of the program code.
- Includes **real world applications** for all foundation models.
- Does not neglect **neuroscience findings** and more neurobiologically oriented models.
- Includes unique and distinctive chapters on **neuroscienc**e, **neural networks and statistical pattern recognition, support vector machines, pulsed neural networks, fuzzy sets and systems, soft computing**, and **dynamical systems**.

Marginal notes (like this one) provide remarks, mathematical tips to understand equations, and occasional historical or bibliographic pointers.

In addition, at the end of each chapter, **Chapter Summaries** provide compact overviews of material covered, and **Bibliographic Remarks** provide directions for researchers. Each Chapter (except 1, 2, 13, 15) ends with a set of exercises that review important ideas developed in the text. Some of these require MATLAB programs to be written. They help develop a clear understanding of the working of an algorithm while strengthening logical thinking.

Organization of the Book

This book is roughly divided into four broad parts.

Part I: History and Neuroscience

Historical developments in the exciting field of neuroscience from antiquity are first reviewed in Chapter 1 to motivate the reader, and to find neural networks a place in the artificial intelligence and soft computing paradigms. This is followed by a detailed review of the global functions of the human brain and the inner working of biological neurons in Chapter 2. Many texts skip such details on neuroscience. However, as is now being increasingly recognized by the international research community, neural network models cannot any longer sacrifice biological fidelity. It is therefore necessary to work through conventional neural network algorithms while relating the underlying theme to cutting edge neuroscience research findings. Such a premise provides the fabric for presentation of ideas throughout this text.

Part II: Supervised Feedforward Systems

The working of artificial neurons and their architectures stands in stark contrast with their biological counterparts. The theory of feedforward systems is put on a firm foundation to demonstrate the computational power and limitations of these systems. Chapter 3 develops a working model of an artificial neuron, and introduces neural network architectures, properties and application domains. Chapter 4 works further on the geometry of artificial neurons and simple networks, to develop an intuitive understanding of fundamental design principles. In Chapter 5 we gently introduce the idea of supervised learning in the context of the Perceptron and LMS models. These concepts are carried to their fullest complexity with backpropagation and reinforcement learning in Chapter 6.

Almost half-way through the text, in Chapter 7, a solid discussion on the relationship between neural networks and statistical pattern recognition is provided. Starting out with Bayes' theorem the discussion works through error functions for feedforward neural networks. This leads naturally to the extremely important topic of generalization which is discussed in the context of statistical learning theory in Chapter 8. Support vector machines and radial basis function networks are given an in-depth treatment. These important topics are missing or have been neglected from many current texts on neural networks.

Part III: Unsupervised Recurrent Systems

In the third part, the book concentrates on feedback neurodynamical models. Chapter 9 lays the foundation for feedback neural network models with an in-depth review on dynamical systems which leads to a study of the models for real neurons and their networks. Current texts miss out on this foundation which is necessary to understand and fully appreciate the beauty and complexity of feedback neural network models. The text continues in Chapter 10 with a treatment of neural networks that are dynamical in nature. Our study takes us through a number of powerful feedback models including attractor neural networks such as the Hopfield network, BSB model, and stochastic models such as the Boltzmann machine. The discussion naturally leads to the bidirectional associative memory model and Hebbian learning.

In Chapter 11, the logical progression of the noise-saturation dilemma, shunting dynamics, instars and outstars, towards the theory of adaptive resonance are introduced. This unified treatment leads to Chapter 12 where the issue of competitive learning is taken up, and the treatment of models culminates in the presentation of the Kohonen self-organizing feature map algorithm.

Part IV: Contemporary Topics

The last part of the textbook covers contemporary topics. Chapter 13 covers the next generation *pulsed neural network* models and provides an in-depth treatment of this recent class of models that take the theory of neural computing full circle back to biology! The discussion touches on the spiking neuron models, integrate-and-fire models, conductance based models and briefly introduces what can be computed with these models. Fuzzy sets and systems are discussed next in Chapter 14, followed by a detailed review of the now ubiquitous soft computing paradigm in Chapter 15. This last chapter outlines applications where neural models have been integrated with other supporting soft computing technologies such as fuzzy systems and evolutionary algorithms. In addition, Appendix A briefly reviews neural network hardware implementations, and Appendix B provides an interesting list of Web pointers on the subject of neural networks and related areas.

Figure 1 shows various trajectories that instructors can take depending on the available time, course objectives, and target audience. The arrows indicate the logical sequences in which various reading trajectories exist. It is recommended that chapters that are clubbed be read in immediate succession.

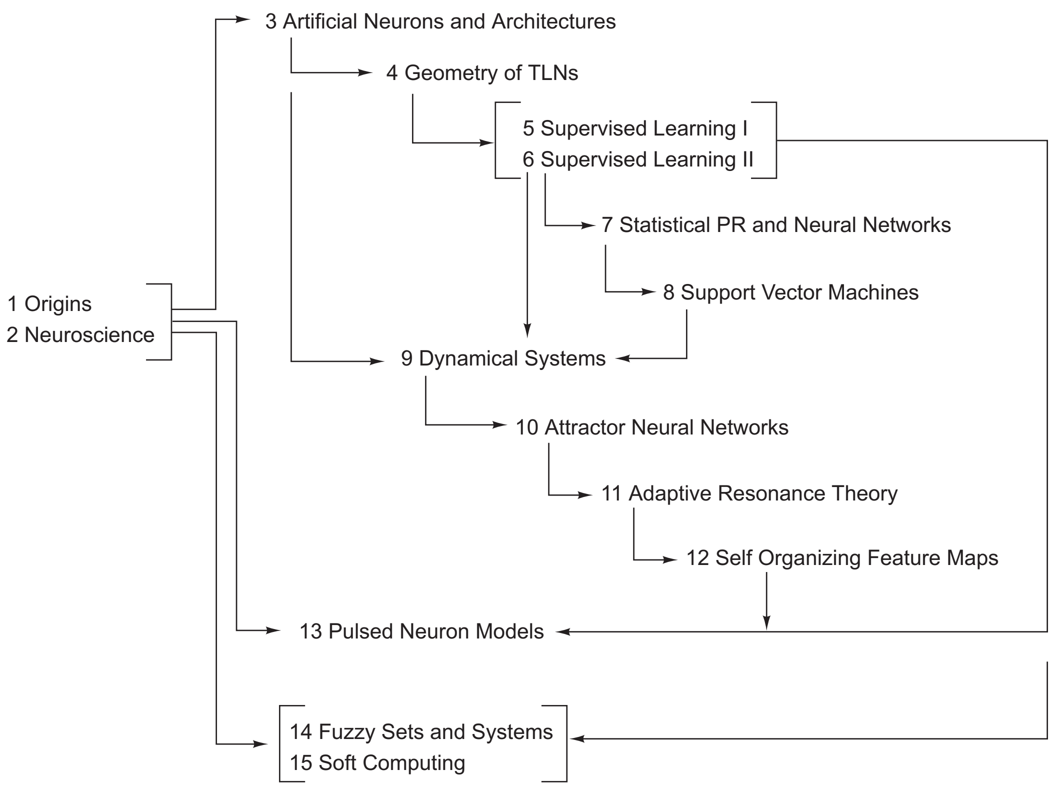

Fig. 1 Chapter dependencies can lead to various possible trajectories.

Road maps for various target courses are suggested in Table 1. They are there to serve as a rough guideline. The ultimate choices are yours.

Target Courses and Chapter Trajectories
Foundation Course on Neural Networks
$1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6^{\dagger} \rightarrow 7 \rightarrow 8 \rightarrow 9 \rightarrow 10^{\dagger\dagger} \rightarrow 11 \rightarrow 12$
Major Course, † Leaving out Reinforcement Learning, †† Leaving out BSB and/or BAM. Chapters 7, 8, 13, 14, 15 as self-reading/seminar.
Introduction to Neural Networks
$1^{\dagger} \rightarrow 2^{\dagger} \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6 \rightarrow 10^{\dagger\dagger}$
Part of AI/Soft Computing Course, † Brief coverage, †† Only Hopfield Network.
Supervised Learning Systems Track
$1^{\dagger} \rightarrow 2^{\dagger} \rightarrow 3^{\dagger\dagger} \rightarrow 4^{\dagger\dagger} \rightarrow 5 \rightarrow 6 \rightarrow 7 \rightarrow 8$
Short Course, † Brief coverage, †† Selecting Essentials.
Unsupervised Learning Systems Track
$1^{\dagger} \rightarrow 2^{\dagger} \rightarrow 3^{\dagger\dagger} \rightarrow 9^{\dagger\dagger} \rightarrow 10 \rightarrow 11 \rightarrow 12$
Short Course, † Brief coverage, †† Selecting Essentials.
Second Level Course on Neural Networks and Soft Computing
$1^{\dagger} \rightarrow 2^{\dagger} \rightarrow 3^{\dagger}, 4^{\dagger}, 5^{\dagger}, 6^{\dagger} \rightarrow 7 \rightarrow 8 \rightarrow 9^{\dagger}, 10^{\dagger}, 11^{\dagger}, 12^{\dagger} \rightarrow 13 \rightarrow 14 \rightarrow 15$
Major Course, † Review, Minor Project Supplement.

Table 1 A few sample road maps for various target courses

Acknowledgements

During the course of writing this book many people have supported me in countless different ways. I express my deepest sense of gratitude to all of them, whether mentioned here or not.

It has been most refreshing to be associated with the dynamic TMH Team. Vikesh Dhyani helped kick-off the project. Vibha Mahajan, Biju Ganesan, Mini Narayanan, and Piyaray Lal Pandita have been great people to work with: full of wonderful ideas, very meticulous, forward looking,

and most patient. Thanks so much to all of you and to everyone else at TMH for bringing this project to its final form.

My sincerest thanks go to Professor Nikhil Pal of the Indian Statistical Institute, Kolkata, for his most insightful suggestions on the first draft of the text, and for so kindly agreeing to write a Foreword for the book. He has been a great source of inspiration. My very good friend Professor Gursaran Adhar of the University of North-Carolina, Wilmington, brimming over with enthusiasm, has been a great source of encouragement: from research reprints to copyright permissions, he has helped throughout.

My own group at the Neural Networks Lab, has demonstrated uncommon and tireless dedication to this project. I cannot thank them enough with only these few words, for their contribution and my gratitude cannot be quantified. Charan Priya tamed the beautiful beast LaTeX and painstakingly typeset the first version of the manuscript. Sandeep and Shunmuga—pillars of the NNLAB—have supported me endlessly in every way. Shubhdeep helped patiently and unhesitatingly with the figures, while proof-reading was supported fully by Nutan and Lotika. Other past members of the Lab: Amboo, Alok, Vikas, and Mahima have also helped with so much.

Dr. Vijai Kumar, Medical Director, Jaslok Hospital, went through various versions of the chapter on neuroscience, and provided immense background support. Professor C.P. Nigam's detailed class notes helped organize my ideas. Professors Arun Sinha, V.G. Das and S.P. Gupta have helped me time and again. Soam Prakash has been a great friend and has stood by me in difficult moments, and Gurpyari helped me with so much more than just English! Thanks also to GP, Anuj, and Sudhish in the US for sending me books on the subject.

Back home, my wife, resilient Ritu—with unstinting support and immense patience—helped push this project through with resolve! My father whose knowledge constantly inspires, and my perfectionist mother and sister who have this misconception that I'm a really learned "Prof!", have stood by throughout. My whirlwind mother-in-law, Inderjit, and Neelu, Gursaran and Rajendra: rock solid support. And last of all, dear Ami and Daya, the simple and true joys of my life, who constantly remind me of what artificial neural networks *can never be*.

I couldn't have done this without you all.

SATISH KUMAR