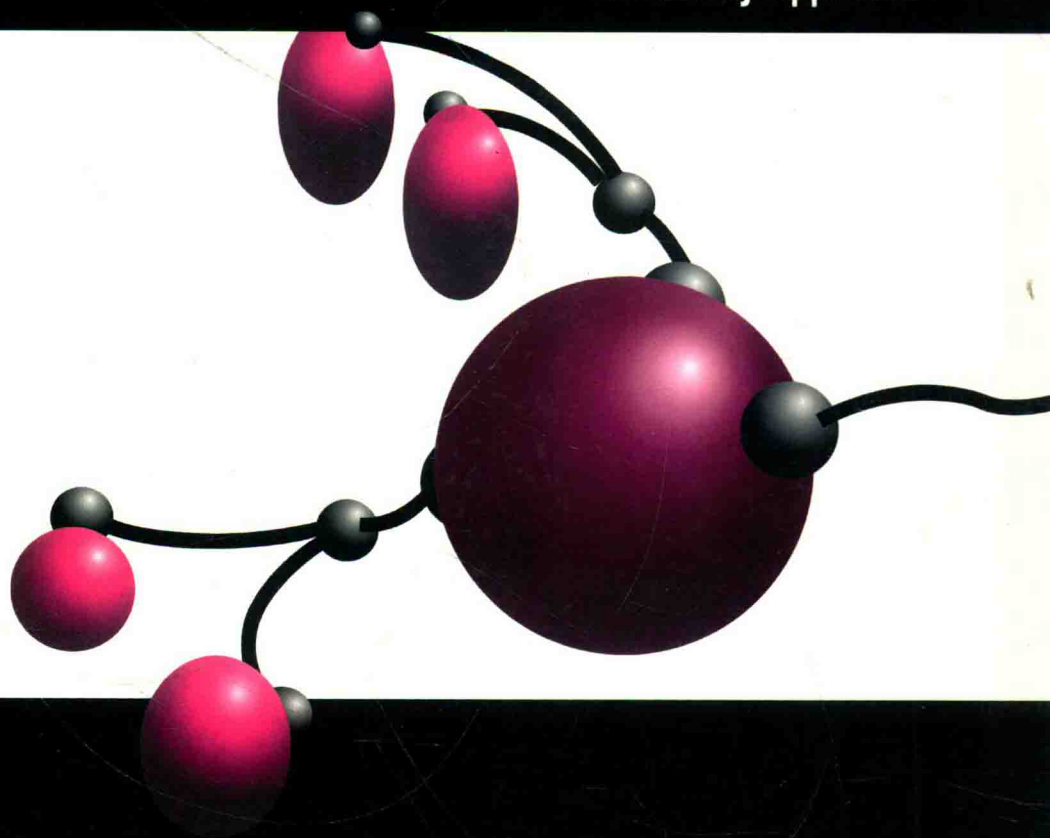


NEURAL NETWORKS AND QUALITATIVE PHYSICS

A Viability Approach



Jean-Pierre Aubin

Neural Networks and Qualitative Physics

JEAN-PIERRE AUBIN

*Ecole Doctorale de Mathématiques de la Décision
Université de Paris - Dauphine*



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town,
Singapore, São Paulo, Delhi, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9781107402843

© Cambridge University Press 1996

This publication is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without the written
permission of Cambridge University Press.

First published 1996
First paperback edition 2011

A catalogue record for this publication is available from the British Library

ISBN 978-0-521-44532-0 Hardback
ISBN 978-1-107-40284-3 Paperback

Cambridge University Press has no responsibility for the persistence or
accuracy of URLs for external or third-party internet websites referred to in
this publication, and does not guarantee that any content on such websites is,
or will remain, accurate or appropriate.

This book is devoted to some mathematical methods that arise in two domains of artificial intelligence: neural networks and qualitative physics. The rapid advances in these two areas have left unanswered several mathematical questions that should motivate and challenge mathematicians.

Professor Aubin makes use of control and viability theory in neural networks and cognitive systems, regarded as dynamical systems controlled by synaptic matrices, and set-valued analysis that plays a natural and crucial role in qualitative analysis and simulation. This allows many examples of neural networks to be presented in a unified way. In addition, several results on the control of linear and nonlinear systems are used to obtain a "learning algorithm" of pattern classification problems, such as the back-propagation formula, as well as learning algorithms of feedback regulation laws of solutions to control systems subject to state constraints.

Mathematical models involve many features of a problem that may not be relevant to its solution. Qualitative physics, however, deals with an imperfect knowledge of the problem model. It is therefore more suited to the study of expert systems, which are shallow models and do not require structural knowledge of the problem.

This book should be a valuable introduction to the field for researchers in neural networks and cognitive systems, and should help to expand the range of study for viability theorists.

Neural Networks and Qualitative Physics

Preface

This book is devoted to some **mathematical methods** that arise in two domains of artificial intelligence: neural networks and qualitative physics (which here we shall call “qualitative analysis”). These two topics are treated independently. Rapid advances in these two areas have left unanswered many mathematical questions that should motivate and challenge a wide range of mathematicians. The mathematical techniques that I choose to present in this book are as follows:

control and viability theory in neural networks and cognitive systems, regarded as dynamical systems controlled by synaptic matrices.
set-valued analysis, which plays a natural and crucial role in qualitative analysis and simulation by emphasizing properties common to a class of problems, data, and solutions. Set-valued analysis also underlies *mathematical morphology*,¹ which provides useful techniques for image recognition.

This allows us to present in a unified way many examples of neural networks and to use several results on the control of linear and nonlinear systems to obtain a *learning algorithm* of pattern-classification problems (including time series in forecasting), such as the *back-propagation formula*, in addition to learning algorithms concerning feedback-regulation laws for solutions to control systems subject to state constraints (inverse dynamics).

These mathematical techniques may also serve to contribute to the various attempts to devise mathematical metaphors for cognitive pro-

¹See the forthcoming book by Michel Schmitt and Luc Vincent, *Morphological Image Analysis* (Cambridge University Press). The links between mathematical morphology and set-valued analysis and viability theory will be explored in a subsequent book, *Mutational and Morphological Analysis: tools for shape regulation and optimization*.

cesses. I present here a very speculative metaphor under the name of *cognitive systems* based on these mathematical techniques. They go beyond neural networks in the sense that they involve the problem of adaptation to viability constraints. They can *recognize* the state of the environment and *act* on the environment to *adapt* to given viability constraints. Instead of encoding knowledge in synaptic matrices as neural networks do, the knowledge is stored in *conceptual controls*. Given the mechanism of recognition of the state of the environment by conceptual controls, perception and action laws, and viability constraints, the viability theorems allow one to construct *learning rules* that describe how conceptual controls evolve in terms of sensorimotor states to adapt to viability constraints.

There is always a combination of two basic motivations for dealing with formal models of cognition - *neural networks* being content with implementation of "neural-like" systems on computers, and *cognitive systems* attempting to model actual biological nervous systems. Every model lies between these two requirements - the first allowing more freedom in the choice of a particular representation (computing efficiency being the main criterion), and the second constraining the modeling to be closer to biological reality.²

The symbolic processing capabilities that neural networks try to achieve are unexpected technological consequences of the use of digital computers, which were not designed for such a purpose at their inception. In the same way, the skills of reasoning logically and solving mathematical problems also represent a kind of unexpected "technological fallout" of the human brain, because they certainly were not among the advantages necessary for the survival of the human species when they appeared.

Expert systems are *shallow models* that do not require any formal and structural knowledge of the problem, whereas a mathematical model might involve too many features that would not be relevant for solving the problem at hand. For many problems, we have only imperfect knowledge of the model, and we may be interested in only a few features (often of a qualitative nature) of the solution, and so we see at once that the concept of *partial knowledge* involves two types of ideas:

² Actually, we should say "degree of reality for a social group at a given time," which is understood here in terms of the consensus interpretations of the group members' perceptions of their physical, biological, social, and cultural environments. This concept of reality is thus relative to a social group and is subject to evolution.

1. We require less precision in the results (e.g., signs of the components of vectors instead of their numerical values),
2. We take into account a broader universality or robustness of these results with respect to uncertainty, disturbances, and lack of precision.

As in numerical analysis, which deals both with approximation of problems in infinite-dimensional spaces by problems in finite-dimensional spaces and with the algorithms for solving such approximated problems, the problems of qualitative analysis arise at two levels: the passage from quantitative analysis to qualitative analysis (which deals with the association of discrete problems with continuous problems) and the algorithms to solve discrete problems.³ In particular, Kuipers's QSIM algorithm for tracking the monotonicity properties of solutions to differential equations is revisited by placing it in a rigorous mathematical framework. This allows us to determine a priori the *landmarks* (i.e., the states at which the monotonicity properties change) instead of discovering them a posteriori by tracking the qualitative evolution of the solutions to the differential equation. These landmarks delineate *qualitative cells*, in which the monotonicity behaviors of the solutions are the same. Once these qualitative cells are computed, the Dordan QSIM algorithm provides the transition laws from one qualitative cell to the others.

This book is divided into 10 chapters. Chapters 1-7 deal with neural networks and some mathematical background needed to treat them (pseudoinverses, tensor products, gradient methods for convex potentials), Chapter 8 deals with cognitive systems, and Chapters 9 and 10 deal with some mathematical questions raised by qualitative physics, in the static and dynamic cases, respectively.

Chapter 1 provides the definitions of neural networks and learning processes (including the perceptron algorithm) and the *heavy learning algorithm*, which allows learning without forgetting.

Chapter 2 deals with some mathematical tools: pseudoinverses of linear operators and tensor products. Indeed, we have to use the specific structure of the space of synaptic matrices as a tensor product to justify mathematically the *connectionist features* of neural networks. Tensor products *explain* the Hebbian nature of many learning algorithms. This is due to the fact that derivatives of a wide class of nonlinear maps de-

³The first aspect has been quite neglected, and it is the one we shall emphasize in this book.

defined on spaces of synaptic matrices are tensor products and also to the fact that the pseudoinverse of a tensor product of linear operators is the tensor product of their pseudoinverses.

Chapter 3 is devoted to the case of linear neural networks, also called *associative memories*. We begin by showing that the *heavy learning algorithm* for neural networks that are affine with respect to the synaptic matrices (but nonlinear with respect to the signals) has a Hebbian character. We proceed with purely linear networks with a single layer or with a finite number or a continuum of layers. The chapter ends with an introduction to associative memories with gates, which are well adapted to compute Boolean and fuzzy Boolean functions.

Chapter 4 is devoted to the proof of the convergence of the gradient method for minimization problems involving a convex criterion with or without constraints. We discuss an application to the Minover algorithm of Mézard that replaces the perceptron algorithm. Many more features of convex analysis could be used in the study of a class of neural networks, but such results would go beyond the scope of this book and the common knowledge of its expected audience.

Chapter 5 adapts these results to the case of nonlinear networks and presents two main types of learning rules. The first class consists of algorithms derived from the gradient method and includes in particular the back-propagation rule. The second class is composed of learning rules based on the Newton method.

Chapter 6 is devoted to the use of neural networks for finding viable solutions to control systems, that is, solutions to control systems that will satisfy given viability (or state) constraints. The purpose of this chapter is to derive learning processes for *regulation feedback* for control problems through neural networks. Two classes of learning rules are presented. The first, called the class of *external learning rules*, is based on the gradient method (of optimization problems involving nonsmooth functions). The second deals with *uniform algorithms*.

In Chapter 7, the *internal-learning algorithm* provides learning rules based on viability theory. Two sections are devoted to a short presentation of the main results of viability theory and its application to the regulation of viable solutions to control systems. Applications to the control of cart-pole problems and other benchmark problems have been designed by N. Seube. This algorithm is applied to stabilization problems.

Chapter 8 goes beyond neural networks as they are usually defined. It proposes a very speculative mathematical model of what is called a

cognitive system. A cognitive system is a dynamical system describing the evolution of sensorimotor states, recognized and controlled by *conceptual controls*, according to perception and action laws, and required to obey some viability constraints. *Adaptive learning processes* associating conceptual controls with sensorimotor states are then obtained by using viability theorems, including the ones that obey the *inertia principle*: *Change the conceptual controls only when the viability of the cognitive system is at stake*. This chapter is oriented toward mathematical metaphors motivated by cognisciences, of which we present a few relevant facts.

Chapter 9 treats the qualitative resolution of static problems described in the form of both equations and inclusions. It proposes a general framework (*confluence frames*) to link quantitative problems with qualitative ones. In particular, sign confluences are thoroughly investigated.

Chapter 10 is devoted to qualitative simulation of differential equations and to a mathematical treatment of Kuipers's QSIM algorithm to track the monotonicity properties of solutions to differential equations. We also consider Dordan's QSIM algorithm, which provides the qualitative cells delineated by the landmarks, and then the transition map associating with each qualitative cell its successor(s). Dordan's QSIM algorithm was designed to study the qualitative behaviors of a class of differential systems, the *replicator systems*, which play important roles in several domains of biology and biochemistry. We consider several examples obtained by using software designed by O. Dordan.

Two appendixes conclude this book. Appendix A provides a survey of convex optimization and set-valued analysis that goes beyond the minimal survey of Chapter 4. Appendix B describes applications of Nicolas Seube's algorithms, presented in Chapters 6 and 7, to the control of autonomous underwater vehicles (AUVs) tracking the trajectory of an exosystem.

Acknowledgements

I thank Alain Bensoussan for asking me to report to INRIA about neural networks and qualitative physics. That triggered a deeper interest in this topic.

This book owes much to Olivier Dordan and Nicolas Seube, who developed most of the material of this book and the computer applications. I have the pleasure to express here my gratitude. I also thank Luc Doyen and Juliette Mattioli for correcting parts of the text and for developing other applications of set-valued analysis and viability theory to visual control and mathematical morphology.

Contents

<i>Preface</i>	<i>page</i> xi
<i>Acknowledgements</i>	xvii
1 Neural Networks: A Control Approach	1
1.1 Neural Networks: A General Presentation	4
1.2 Examples of Neural Networks	8
1.3 Rosenblatt's Perceptron and the Perceptron Algorithm	13
1.4 Multilayer Neural Networks	15
1.5 Adaptive Systems	18
2 Pseudoinverses and Tensor Products	23
2.1 Pseudoinverses	24
2.2 The Heavy Algorithm Theorem	32
2.3 Duals of Spaces of Linear Operators	34
2.4 Tensor Products of Linear Operators	36
2.5 Pseudoinverses of Tensor Products	38
3 Associative Memories	44
3.1 Optimal Synaptic Matrices of Associative Memories	44
3.2 The Heavy Learning Algorithm for Neural Networks	49
3.3 Multilayer Associative Memories	51
3.4 Continuous-Layer Associative Memories	54
3.5 Associative Memories with Gates	56
4 The Gradient Method	61
4.1 Extended Functions and Their Epigraphs	63
4.2 Convex Functions	64
4.3 Subgradient Algorithm	66
4.4 Support Functions and Conjugate Functions	68
4.5 The Minover Algorithm	73

5	Nonlinear Neural Networks	75
5.1	Gradient Learning Rules	79
5.2	Back-Propagation Algorithm	82
5.3	Continuous Back-Propagation Algorithm	88
5.4	Newton-type Methods	93
6	External Learning Algorithm for Feedback Controls	100
6.1	Construction of the External Learning Algorithm	102
6.2	Simulation Results: The Cart Pole	105
6.3	Simulation Results: A Target Problem	109
6.4	Uniform External Algorithm	113
7	Internal Learning Algorithm for Feedback Controls	118
7.1	Viability Theory at a Glance	119
7.2	Regulation of Control Systems	125
7.3	Internal Learning Algorithm	133
7.4	Stabilization Problems	137
8	Learning Processes of Cognitive Systems	139
8.1	Viable Learning Processes	143
8.2	Motivations and Comments	153
8.3	Recognition Mechanisms	158
9	Qualitative Analysis of Static Problems	160
9.1	Qualitative Frames	164
9.2	Qualitative Equations and Inclusions	167
9.3	Qualitative Operations	171
9.4	Qualitative Solutions to Linear Problems	175
9.5	The Constrained Inverse-Function Theorem	179
9.6	Single-valued Nonlinear Equations	180
9.7	Set-valued Problems	183
10	Dynamical Qualitative Simulation	187
10.1	Monotonic Cells	189
10.2	Transitions between Qualitative Cells	192
10.3	Qualitative Equilibrium and Repeller	195
10.4	The QSIM Algorithm	196
10.5	Replicator Systems	198
10.6	Qualitative Simulation of Replicator Systems	203
10.7	Nonemptiness and Singularity of Qualitative Cells	207
10.8	General Qualitative Cells	209
10.9	Sufficient Conditions for Chaos	217
	<i>Appendix 1</i> Convex and Nonsmooth Analysis	219

Contents

ix

<i>Appendix 2</i>	Control of an AUV	252
<i>Bibliography</i>		262
<i>Index</i>		280

1

Neural Networks: A Control Approach

Introduction

A *neural network* is a network of subunits, called “formal neurons,” processing input signals to output signals, which are coupled through “synapses.” The synapses are the nodes of this particular kind of network, the “strength” of which, called the *synaptic weight*, codes the “knowledge” of the network and controls the processing of the signals.

Let us be clear at the outset that the resemblance of a formal neuron to an animal-brain neuron is not well established, but that is not essential at this stage of abstraction. However, this terminology can be justified to some extent, and it is by now widely accepted, as discussed later. Chapter 8 develops this issue.

Also, there is always a combination of two basic motivations for dealing with neural networks - one attempting to model actual biological nervous systems, the other being content with implementation of neural-like systems on computers. Every model lies between these two requirements - the first constraining the modeling, the second allowing more freedom in the choice of a particular representation.

There are so many different versions of neural networks that it is difficult to find a common framework to unify all of them at a rather concrete level. But one can regard neural networks as dynamical systems (discrete or continuous), the states of which are the signals, and the controls of which are the synaptic weights, which regulate the flux of transmitters from one neuron to another. They yield what are also known as *adaptive systems*, controlled by *synaptic matrices*.

We investigate in this chapter the *supervised learning* of a finite set of patterns, called the *training set*, each pattern being a pair of input-output signals. A *learning process* amounts to matching the given input

signals of the patterns of this training set with the associated output signals. It thus involves a comparison with desired answers, done by a “supervisor.” Hence, the learning problem amounts to finding a synaptic matrix that can “learn” the patterns of a given training set. Any algorithm yielding such a synaptic matrix will then be regarded as a *learning rule*. This seems to exclude neural networks from reasonable models of basic neural functions. However, the problems under investigation here are as follows:

1. Establish the existence of exact (or approximate) synaptic matrices that have learned a given set of patterns,
2. Find algorithms, regarded as learning rules, that will provide a sequence of synaptic matrices converging to a solution of the learning problem.

I propose in this first chapter a short presentation of the learning processes of neural networks in the framework of dynamical systems controlled by synaptic matrices. Section 1.2 explains how neural networks operate to solve pattern-classification problems and, in particular, to extrapolate time series in forecasting problems. These problems are called supervised learning problems, because the patterns to be taught are provided by a supervisor, so to speak.

Biological Comments. Although the results to be derived here will be “formal” and will not necessarily be associated with any biological implementation, it may be useful to provide a crude description of the nature of the processing carried out by some biological neurons (Figure 1.1).

It should be first pointed out that many kinds of neuronal cells evolved during phylogenesis, each of them selected to provide adequate technical solutions to biological or environmental problems. Most neurons in the central nervous system in higher animals can be regarded as *impulse oscillators*. They produce trains or volleys of neural impulses whose average frequency will depend on the input excitation. The processing carried out by neuron ensues from biophysical and biochemical phenomena in the membrane of the neuron, wherein functions the machinery that controls the interaction of the cell and its environment.

Models of neural dynamics that attempt to describe the triggering phenomena, the transmission of the impulses, and their biochemical control have been extensively studied since the Hodgkin-Huxley model was first proposed. The central features are the following: There is a static

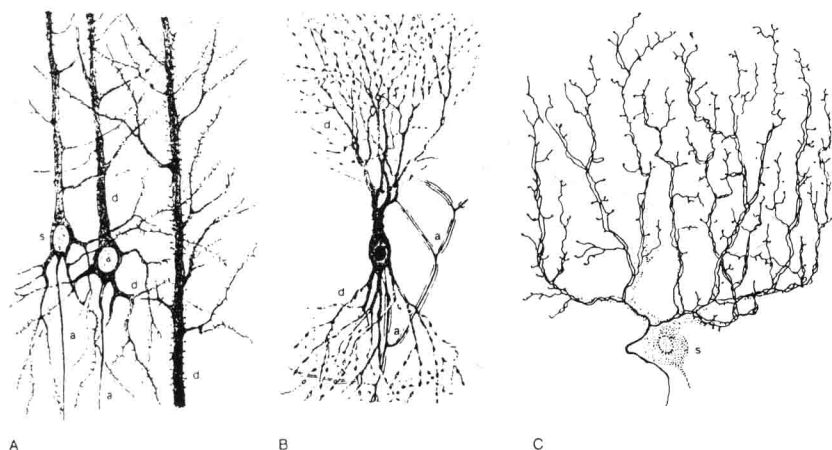


Fig. 1.1. Neurons. Pyramidal and Purkinje neurons

electrical-potential difference of about 70 mV between the inside of the neuron (negatively charged) and its environment; it is maintained by diffusion of ions through the membrane. The synaptic receptors are able to control the ionic conductance of the membrane through a highly sophisticated biochemical and ionic mechanism (Figure 1.2).

When neurotransmitters (excitation signals) arrive at a receptor, the excitatory synapses tend to produce a depolarization of the membrane, and the inhibitory synapses tend to produce a hyperpolarization. When the sum (i.e., algebraic sum) of the depolarizations exceeds a given threshold, then the membrane's permeability to ions is increased, and the membrane becomes electrically active, sending an output impulse of about 100 mV amplitude during a period of 0.5-2 ms. After each impulse, there is a short refractory period during which the membrane recovers so as to be ready for the next impulse (Figure 1.3).

Because in our framework we are interested in the collective processing by neurons, we shall be content with only a crude analytical description of the processing role of the neuron. We retain only the following features: The impulse frequency oscillates between bounds determined by physical and chemical factors (oscillating between 0 and 500 Hz). Between those bounds, the postsynaptic *average oscillatory frequency* is assumed to be a monotonic function of the net algebraic sum of the presynaptic average frequencies of the inputs afferent to the neurons (inhibitory excitations being regarded as negative excitations). Even though, for simplicity, the integration of the presynaptic inputs is often