

Stochastic Mechanics
Random Media
Signal Processing
and Image Synthesis
Mathematical Economics
Stochastic Optimization
Stochastic Control

**Applications of
Mathematics**
*Stochastic Modelling
and Applied Probability*

35

Edited by I. Karatzas
M. Yor

Advisory Board P. Brémaud
E. Carlen
R. Dobrushin
W. Fleming
D. Geman
G. Grimmett
G. Papanicolaou
J. Scheinkman

Applications of Mathematics

- 1 Fleming/Rishel, **Deterministic and Stochastic Optimal Control** (1975)
- 2 Marchuk, **Methods of Numerical Mathematics**, Second Ed. (1982)
- 3 Balakrishnan, **Applied Functional Analysis**, Second Ed. (1981)
- 4 Borovkov, **Stochastic Processes in Queueing Theory** (1976)
- 5 Liptser/Shiryayev, **Statistics of Random Processes I: General Theory** (1977)
- 6 Liptser/Shiryayev, **Statistics of Random Processes II: Applications** (1978)
- 7 Vorob'ev, **Game Theory: Lectures for Economists and Systems Scientists** (1977)
- 8 Shiryayev, **Optimal Stopping Rules** (1978)
- 9 Ibragimov/Rozanov, **Gaussian Random Processes** (1978)
- 10 Wonham, **Linear Multivariable Control: A Geometric Approach**, Third Ed. (1985)
- 11 Hida, **Brownian Motion** (1980)
- 12 Hestenes, **Conjugate Direction Methods in Optimization** (1980)
- 13 Kallianpur, **Stochastic Filtering Theory** (1980)
- 14 Krylov, **Controlled Diffusion Processes** (1980)
- 15 Prabhu, **Stochastic Storage Processes: Queues, Insurance Risk, and Dams** (1980)
- 16 Ibragimov/Has'minskii, **Statistical Estimation: Asymptotic Theory** (1981)
- 17 Cesari, **Optimization: Theory and Applications** (1982)
- 18 Elliott, **Stochastic Calculus and Applications** (1982)
- 19 Marchuk/Shaidourov, **Difference Methods and Their Extrapolations** (1983)
- 20 Hijab, **Stabilization of Control Systems** (1986)
- 21 Protter, **Stochastic Integration and Differential Equations** (1990)
- 22 Benveniste/Métivier/Priouret, **Adaptive Algorithms and Stochastic Approximations** (1990)
- 23 Kloeden/Platen, **Numerical Solution of Stochastic Differential Equations** (1992)
- 24 Kushner/Dupuis, **Numerical Methods for Stochastic Control Problems in Continuous Time** (1992)
- 25 Fleming/Soner, **Controlled Markov Processes and Viscosity Solutions** (1993)
- 26 Baccelli/Brémaud, **Elements of Queueing Theory** (1994)
- 27 Winkler, **Image Analysis, Random Fields, and Dynamic Monte Carlo Methods: An Introduction to Mathematical Aspects** (1994)
- 28 Kalpazidou, **Cycle Representations of Markov Processes** (1995)
- 29 Elliott/Aggoun/Moore, **Hidden Markov Models: Estimation and Control** (1995)
- 30 Hernández-Lerma/Lasserre, **Discrete-Time Markov Control Processes: Basic Optimality Criteria** (1996)
- 31 Devroye/Györfi/Lugosi, **A Probabilistic Theory of Pattern Recognition** (1996)
- 32 Maitra/Sudderth, **Discrete Gambling and Stochastic Games** (1996)
- 33 Embrechts/Klüppelberg/Mikosch, **Modelling Extremal Events** (1997)
- 34 Duflo, **Random Iterative Models** (1997)
- 35 Kushner/Yin, **Stochastic Approximation Algorithms and Applications** (1997)
- 36 Musiela/Rutkowski, **Martingale Methods in Financial Modeling: Theory and Application** (1997)

Harold J. Kushner G. George Yin

Stochastic Approximation Algorithms and Applications

With 24 Figures



Springer

Harold J. Kushner
Division of Applied Mathematics
Brown University
Providence, RI 02912, USA

G. George Yin
Department of Mathematics
Wayne State University
Detroit, MI 48202, USA

Managing Editors

I. Karatzas
Departments of Mathematics and Statistics
Columbia University
New York, NY 10027, USA

M. Yor
CNRS, Laboratoire de Probabilités
Université Pierre et Marie Curie
4, Place Jussieu, Tour 56
F-75252 Paris Cedex 05, France

Mathematics Subject Classification (1991): 62L20, 93E10, 93E23, 65C05, 93E35, 93-02, 90C15

Library of Congress Cataloging-in-Publication Data

Kushner, Harold J. (Harold Joseph), 1933–

Stochastic approximation algorithms and applications/Harold J.

Kushner, G. George Yin.

p. cm. — (Applications of mathematics; 35)

Includes bibliographical references and index.

ISBN 978-1-4899-2698-2

1. Stochastic approximation. I. Yin, George, 1954–

II. Title. III. Series.

QA274.2.K88 1997

519.5'4—dc21

96-48847

Printed on acid-free paper.

ISBN 978-1-4899-2698-2 ISBN 978-1-4899-2696-8 (eBook)

DOI 10.1007/978-1-4899-2696-8

© 1997 Springer Science+Business Media New York

Originally Published by Springer-Verlag New York, Inc in 1997

Softcover reprint of the hardcover 1st edition 1997

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher, except for brief excerpts in connection with reviews or scholarly analysis.

Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Production managed by Francine McNeill; manufacturing supervised by Jacqui Ashri.
Photocomposed copy prepared from the authors' LaTeX files.

9 8 7 6 5 4 3 2 1

To
Linda, Diana and Nina
and
Chen

Contents

Preface and Introduction	xiii
1 Introduction: Applications and Issues	1
1.0 Outline of Chapter	1
1.1 The Robbins–Monro Algorithm	2
1.1.1 Introduction	2
1.1.2 Finding the Zeros of an Unknown Function	5
1.1.3 A Linear Pattern Classifier: Best Linear Least Squares Fit	8
1.1.4 Minimization by Recursive Monte Carlo	11
1.2 The Kiefer–Wolfowitz Procedure	13
1.2.1 The Basic Procedure	13
1.2.2 Random Directions	16
1.3 Extensions of the Algorithms: Variance Reduction, Robustness, Iterate Averaging, Constraints, and Convex Optimization	18
1.3.1 A Variance Reduction Method	18
1.3.2 Constraints	20
1.3.3 Averaging of the Iterates: “Polyak Averaging”	21
1.3.4 Robust Algorithms	22
1.3.5 Nonexistence of the Derivative at Some θ	22
2 Applications to Learning, State Dependent Noise, and Queueing	25
2.0 Outline of Chapter	25

2.1	An Animal Learning Model	26
2.2	A Neural Network	29
2.3	Q -Learning	32
2.4	State Dependent Noise: A Motivational Example	35
2.5	Optimization of a GI/G/1 Queue	38
2.5.1	Derivative Estimation and Infinitesimal Perturbation Analysis: A Brief Review	39
2.5.2	The Derivative Estimate for the Queueing Problem	41
2.6	Passive Stochastic Approximation	45
3	Applications in Signal Processing and Adaptive Control	47
3.0	Outline of Chapter	47
3.1	Parameter Identification and Tracking	48
3.1.1	The Classical Model	48
3.1.2	ARMA and ARMAX Models	51
3.2	Tracking Time Varying Systems: An Adaptive Step Size Algorithm	53
3.2.1	The Algorithm	53
3.2.2	Some Data	56
3.3	Feedback and Averaging in the Identification Algorithm	58
3.4	Applications in Communication Theory	60
3.4.1	Adaptive Noise Cancellation and Disturbance Rejection	60
3.4.2	Adaptive Equalizers	62
4	Mathematical Background	67
4.0	Outline of Chapter	67
4.1	Martingales, Submartingales, and Inequalities	68
4.2	Ordinary Differential Equations	72
4.2.1	Limits of a Sequence of Continuous Functions	72
4.2.2	Stability of Ordinary Differential Equations	74
4.3	Projected ODE	77
4.4	Stochastic Stability and Perturbed Stochastic Liapunov Functions	80
5	Convergence with Probability One: Martingale Difference Noise	85
5.0	Outline of Chapter	85
5.1	Truncated Algorithms: Introduction	87
5.2	The ODE Method: A Basic Convergence Theorem	93
5.2.1	Assumptions and the Main Convergence Theorem	93
5.2.2	Chain Recurrence	102
5.3	A General Compactness Method	107
5.3.1	The Basic Convergence Theorem	107

5.3.2	Sufficient Conditions for the Rate of Change Condition	109
5.3.3	The Kiefer–Wolfowitz Algorithm	113
5.4	Stability and Stability–ODE Methods	114
5.5	Soft Constraints	120
5.6	Random Directions, Subgradients, and Differential Inclusions	122
5.7	Convergence for the Lizard Learning and Pattern Classification Problems	125
5.7.1	The Lizard Learning Problem	125
5.7.2	The Pattern Classification Problem	126
5.8	Convergence to a Local Minimum: A Perturbation Method	127
6	Convergence with Probability One: Correlated Noise	135
6.0	Outline of Chapter	135
6.1	A General Compactness Method	136
6.1.1	Introduction and General Assumptions	136
6.1.2	The Basic Convergence Theorem	140
6.1.3	Local Convergence Results	143
6.2	Sufficient Conditions for the Rate of Change Assumptions: Laws of Large Numbers	144
6.3	Perturbed State Criteria for the Rate of Change Assumptions	146
6.3.1	Introduction to Perturbed Test Functions	146
6.3.2	General Conditions for the Asymptotic Rate of Change	148
6.3.3	Alternative Perturbations	151
6.4	Examples Using State Perturbation	154
6.5	Kiefer–Wolfowitz Algorithms	157
6.6	A State Perturbation Method and State Dependent Noise	159
6.7	Stability Methods	162
6.8	Differential Inclusions and the Parameter Identification Problem	167
6.9	State Perturbation–Large Deviations Methods	168
6.10	Large Deviations Estimates	173
6.10.1	Two-Sided Estimates	173
6.10.2	Upper Bounds and Weaker Conditions	179
6.10.3	Escape Times	182
7	Weak Convergence: Introduction	185
7.0	Outline of Chapter	185
7.1	Introduction	186
7.2	Martingale Difference Noise	189

7.3	Weak Convergence	198
7.3.1	Definitions	198
7.3.2	Basic Convergence Theorems	201
7.4	Martingale Limit Processes and the Wiener Process	205
7.4.1	Verifying that a Process Is a Martingale	205
7.4.2	The Wiener Process	207
7.4.3	A Perturbed Test Function Method for Verifying Tightness and the Wiener Process	208
8	Weak Convergence Methods for General Algorithms	213
8.0	Outline of Chapter	213
8.1	Assumptions: Exogenous Noise and Constant Step Size . .	215
8.2	Convergence: Exogenous Noise	218
8.2.1	Constant Step Size: Martingale Difference Noise .	218
8.2.2	Correlated Noise	225
8.2.3	Step Size $\epsilon_n \rightarrow 0$	228
8.2.4	Random ϵ_n	231
8.2.5	Differential Inclusions	231
8.3	The Kiefer–Wolfowitz Algorithm	232
8.3.1	Martingale Difference Noise	232
8.3.2	Correlated Noise	234
8.4	Markov State Dependent Noise	238
8.4.1	Constant Step Size	238
8.4.2	Decreasing Step Size $\epsilon_n \rightarrow 0$	242
8.4.3	The Invariant Measure Method: Constant Step Size	244
8.4.4	An Alternative Form	246
8.5	Unconstrained Algorithms	247
9	Applications: Proofs of Convergence	251
9.0	Outline of Chapter	251
9.1	Average Cost per Unit Time Criteria: Introduction	252
9.1.1	General Comments	252
9.1.2	A Simple Illustrative SDE Example	254
9.2	A Continuous Time Stochastic Differential Equation Example	258
9.3	A Discrete Example: A GI/G/1 Queue	263
9.4	Signal Processing Problems	266
10	Rate of Convergence	273
10.0	Outline of Chapter	273
10.1	Exogenous Noise: Constant Step Size	274
10.1.1	Martingale Difference Noise	275
10.1.2	Correlated Noise	283
10.2	Exogenous Noise: Decreasing Step Size	286

10.3	The Kiefer–Wolfowitz Algorithm	290
10.3.1	Martingale Difference Noise	290
10.3.2	Correlated Noise	294
10.4	Tightness of the Normalized Iterates: Decreasing Step Size, W.P.1 Convergence	298
10.4.1	Martingale Difference Noise: Robbins–Monro Algorithm	298
10.4.2	Correlated Noise	301
10.4.3	The Kiefer–Wolfowitz Algorithm	303
10.5	Tightness of the Normalized Iterates: Weak Convergence .	305
10.5.1	The Unconstrained Algorithm	305
10.5.2	The Constrained Algorithm and Local Methods . .	307
10.6	Weak Convergence to a Wiener Process	310
10.7	Random Directions: Martingale Difference Noise	315
10.7.1	Comparison of Algorithms	317
10.8	State Dependent Noise	321
11	Averaging of the Iterates	327
11.0	Outline of Chapter	327
11.1	Rate of Convergence of the Averaged Iterates: Minimal Window of Averaging	330
11.1.1	The Robbins–Monro Algorithm: Decreasing Step Size	330
11.1.2	Constant Step Size	333
11.1.3	The Kiefer–Wolfowitz Algorithm	333
11.2	A Two Time Scale Interpretation	335
11.3	Maximal Window of Averaging	336
11.4	The Parameter Identification Problem: An Optimal Algorithm	344
12	Distributed/Decentralized and Asynchronous Algorithms	347
12.0	Outline of Chapter	347
12.1	Examples	349
12.1.1	Introductory Comments	349
12.1.2	Pipelined Computations	350
12.1.3	A Distributed and Decentralized Network Model .	352
12.1.4	Multiaccess Communications	354
12.2	Introduction: Real-Time Scale	355
12.3	The Basic Algorithms	360
12.3.1	Constant Step Size: Introduction	360
12.3.2	Martingale Difference Noise	362
12.3.3	Correlated Noise	368
12.3.4	Infinite Time Analysis	370
12.4	Decreasing Step Size	373
12.5	State Dependent Noise	378

12.6	Rate of Convergence: The Limit Rate Equations	381
12.7	Stability and Tightness of the Normalized Iterates	386
12.7.1	The Unconstrained Algorithm	386
12.8	Convergence for Q -Learning: Discounted Cost	390
References		393
Symbol Index		409
Index		413

Preface and Introduction

The basic stochastic approximation algorithms introduced by Robbins and Monro and by Kiefer and Wolfowitz in the early 1950s have been the subject of an enormous literature, both theoretical and applied. This is due to the large number of applications and the interesting theoretical issues in the analysis of “dynamically defined” stochastic processes. The basic paradigm is a stochastic difference equation such as $\theta_{n+1} = \theta_n + \epsilon_n Y_n$, where θ_n takes its values in some Euclidean space, Y_n is a random variable, and the “step size” $\epsilon_n > 0$ is small and might go to zero as $n \rightarrow \infty$. Roughly speaking, θ is a parameter of a system, and the random vector Y_n is a function of “noise-corrupted” observations taken on the system when the parameter is set to θ_n . One recursively adjusts the parameter so that some goal is met asymptotically. This book is concerned with the qualitative and asymptotic properties of such recursive algorithms in the diverse forms in which they arise in applications. There are analogous continuous time algorithms, but the conditions and proofs are generally very close to those for the discrete time case.

The original work was motivated by the problem of finding a root of a continuous function $\bar{g}(\theta)$, where the function is not known but the experimenter is able to take “noisy” measurements at any desired value of θ . Recursive methods for root finding are common in classical numerical analysis, and it is reasonable to expect that appropriate stochastic analogs would also perform well.

In one classical example, θ is the level of dosage of a drug, and the function $\bar{g}(\theta)$ (which is assumed to be increasing with θ) is the probability of success at dosage level θ . The level at which $\bar{g}(\theta)$ takes a given value

v is sought. The probability of success is known only by experiment at whatever values of θ are selected by the experimenter, with the experimental outcome being either success or failure. Thus, the problem cannot be solved analytically. One possible approach is to take a sufficient number of observations at some fixed value of θ , so that a good estimate of the function value is available, and then to move on. Since most such observations will be taken at parameter values far from the optimum, much effort might be wasted in comparison with the stochastic approximation algorithm $\theta_{n+1} = \theta_n + \epsilon_n[v - \text{observation at } \theta_n]$, where the parameter value moves (on the average) in the correct direction after each observation. In another example, we wish to minimize a real-valued continuously differentiable function $f(\cdot)$ of θ . Here, θ_n is the n th estimate of the minimum, and Y_n is a noisy estimate of the negative of the derivative of $f(\cdot)$ at θ_n , perhaps obtained by a Monte Carlo procedure. The algorithms are frequently constrained in that the iterates θ_n are projected back to some set H if they ever leave it. The mathematical paradigms have posed substantial challenges in the asymptotic analysis of recursively defined stochastic processes.

A major insight of Robbins and Monro was that, if the step sizes in the parameter updates are allowed to go to zero in an appropriate way as $n \rightarrow \infty$, then there is an implicit averaging that eliminates the effects of the noise in the long run. An excellent survey of developments up to about the mid 1960s can be found in the book by Wasan [181]. More recent material can be found in [9, 34, 42, 99, 165]. The book [139] deals with many of the issues involved in stochastic optimization in general.

In recent years, algorithms of the stochastic approximation type have found applications in new and diverse areas, and new techniques have been developed for proofs of convergence and rate of convergence. The actual and potential applications in signal processing have exploded. Indeed, whether or not they are called stochastic approximations, such algorithms occur frequently in practical systems for the purposes of noise or interference cancellation, the optimization of “post processing” or “equalization” filters in time varying communication channels, adaptive antenna systems, and many related applications. In these applications, the step size is often a small constant $\epsilon_n = \epsilon$, or it might be random. In such applications, the underlying processes are often nonstationary, the optimal value of θ changes with time, and one keeps ϵ_n strictly away from zero in order to allow “tracking.” Such tracking applications lead to new problems in the asymptotic analysis (e.g., when ϵ_n are adjusted adaptively); one wishes to estimate the tracking errors and their dependence on the structure of the algorithm.

New challenges have arisen in applications to adaptive control. There has been a resurgence of interest in general “learning” algorithms, motivated by the training problem in artificial neural networks [4, 29, 70], the on-line learning of optimal strategies in very high-dimensional Markov decision problems [183] with unknown transition probabilities, and in learning

automata [113] and related areas. The actual recursive forms of the algorithms in many such applications are of the stochastic approximation type. Owing to the types of averaging methods used, the “noise” can be “pseudorandom” [134], rather than random.

Owing to the recent extensive development of methods such as infinitesimal perturbation analysis [73] for the estimation of the pathwise derivatives of complex discrete event systems, the possibilities for the recursive on-line optimization of many such systems that arise in communications or manufacturing have been widely recognized. The appropriate algorithms are often of the stochastic approximation type, but the criterion to be minimized is the average cost per unit time over the infinite time interval.

Iterate averaging methods [107, 142, 195], which yield nearly optimal algorithms under broad conditions, have been developed. The iterate averaging effectively adds an additional time scale to the algorithm. Decentralized or asynchronous algorithms introduce new difficulties for analysis. Consider, for example, a problem where computation is split among several processors, operating and transmitting data to one another asynchronously. Such algorithms are only beginning to come into prominence, due to both the developments of decentralized processing and applications where each of several locations might control or adjust “local variables” but the criterion of concern is global. Another current “decentralized” application is in Q -learning, where the component updated at any time depends on the state of a Markov process.

Despite their successes, the classical methods are not adequate for many of the algorithms that arise in such applications. Some of the reasons concern the greater flexibility desired for the step sizes, more complicated dependence properties of the noise and iterate processes, the types of constraints that might occur, ergodic cost functions, possibly additional time scales, nonstationarity and issues of tracking for time varying systems, data-flow problems in the decentralized algorithm, iterate averaging algorithms, desired stronger rate of convergence results, and so forth.

Much modern analysis of the algorithms uses the so-called ODE (ordinary differential equation) method introduced by Ljung [119] and extensively developed by Kushner and coworkers [89, 99, 102] to cover quite general noise processes and constraints by the use of weak ergodic or averaging conditions. The main idea is to show that asymptotically the noise effects average out so that the asymptotic behavior is determined effectively by that of a “mean” ODE. The usefulness of the technique stems from the fact that the ODE is obtained by a “local analysis,” where the dynamical term of the ODE at parameter value θ is obtained by averaging the Y_n as though the parameter were fixed at θ . Constraints, complicated state dependent noise processes, discontinuities, and many other difficulties can be handled. Depending on the application, the ODE might be replaced by a constrained (projected) ODE or a differential inclusion. Owing to its versatility and naturalness, the ODE method has become a fundamental

technique in the current toolbox, and its full power will be apparent from the results in this book.

The first three chapters describe applications and serve to motivate the assumptions and theorems to follow. Chapter 1 provides the general motivation underlying stochastic approximation and describes various classical examples. Modifications of the algorithms due to robustness concerns, improvements based on iterate averaging methods, variance reduction, and other modeling issues are also introduced. Chapter 2 contains more advanced examples, each of which is typical of a large class of current interest: animal adaptation models, the so-called Q -learning, and artificial neural networks. The optimization of discrete event systems is introduced by the application of infinitesimal perturbation analysis to the optimization of the performance of a queue with an ergodic cost criterion. The mathematical and modeling issues raised in this example are typical of many of the optimization problems in discrete event systems or where ergodic cost criteria are involved. The example also serves to introduce the concept of “state dependent noise.” Chapter 3 describes some applications arising in adaptive control and communication theory, areas that are major users of stochastic approximation algorithms. Some of the mathematical results that will be needed are collected in Chapter 4.

The book gives quite general combined “stability–ODE” methods for unconstrained problems. Nevertheless, a large part of the work concerns constrained algorithms, because constraints are generally present either explicitly or implicitly. For example, in the queue optimization problem of Chapter 2, the parameter to be selected controls the service rate. What is to be done if the service rate at some iteration is considerably larger than any possible practical value? Either there is a problem with the model or the chosen step sizes, or some bizarre random numbers appeared. Furthermore, in practice the “physics” of models at large parameter values are often poorly known or inconvenient to model, so that whatever “convenient mathematical assumptions” made might be meaningless. No matter what the cause is, one would normally alter the unconstrained algorithm if the parameter θ took on excessive values. The simplest alteration is truncation. Of course, in addition to truncation, a practical algorithm would have other safeguards to ensure robustness against “bad” noise or inappropriate step sizes, etc. It has been somewhat traditional to allow the iterates to be unbounded and to use stability methods to prove that they do, in fact, converge. This approach still has its place and is dealt with here. Indeed, one might even alter the dynamics by introducing “soft” constraints, which have the desired stabilizing effect. However, allowing unbounded iterates seems to be of greater mathematical than practical interest. Owing to the interest in the constrained algorithm, the “constrained ODE” is also discussed in Chapter 4. The chapter contains a brief discussion of stochastic stability and the perturbed stochastic Liapunov function, which play an essential role in the asymptotic analysis.

The first convergence results appear in Chapter 5, which deals with the classical case where the Y_n can be written as the sum of a conditional mean $g_n(\theta_n)$ and a noise term, which is a “martingale difference.” The basic techniques of the ODE method are introduced, both with and without constraints. It is shown that, under reasonable conditions on the noise, there will be convergence with probability one to a “stationary point” or “limit trajectory” of the mean ODE for step size sequences that decrease at least as fast as $\alpha_n/\log n$, where $\alpha_n \rightarrow 0$. If the limit trajectory of the ODE is not concentrated at a single point, then the asymptotic path of the stochastic approximation is concentrated on a limit or invariant set of the ODE that is also “chain recurrent” [6, 62]. Equality constrained problems are included in the basic setup.

Much of the analysis works with interpolated processes. The iterates θ_n are interpolated into a continuous time process with interpolation intervals ϵ_n . The asymptotics (large n) of the iterate sequence are also the asymptotics (large t) of this interpolated sequence. It is the paths of the interpolated process that are approximated by the paths of the ODE.

If there are no constraints, a stability method is used to show that the iterate sequence is recurrent. From this point on, the proofs are a special case of those for the constrained problem. As an illustration of the methods, convergence is proved for an animal learning example (where the step sizes are random, depending on the actual history) and a pattern classification problem. In the minimization of convex functions, the subdifferential replaces the derivative, and the ODE becomes a differential inclusion, but the convergence proofs carry over.

Chapter 6 treats probability one convergence with correlated noise sequences. The development is based on the general “compactness methods” of [99]. The assumptions on the noise sequence are implied by (but weaker than) strong laws of large numbers. In some cases, they are both necessary and sufficient for convergence. They are intuitively natural, and probably not far from being necessary and sufficient even more generally. The way the conditions are formulated allows us to use simple and classical compactness methods to define the mean ODE and to show that its asymptotics characterize that of the algorithm. Stability methods for the unconstrained problem and the generalization of the ODE to a differential inclusion are discussed. The methods of large deviations theory provide an alternative approach to proving convergence under weak conditions, and some simple results are presented.

In Chapters 7 and 8, we work with another type of convergence, called *weak convergence*, since it is based on the theory of weak convergence of a sequence of probability measures and is weaker than convergence with probability one. It is actually much easier to use in that convergence can be proved under weaker and more easily verifiable conditions and generally with substantially less effort. The approach yields virtually the same information on the asymptotic behavior. The weak convergence methods

have considerable theoretical and modeling advantages when dealing with complex problems involving correlated noise, state dependent noise, decentralized or asynchronous algorithms, and discontinuities in the algorithm. It will be seen that the conditions are often close to minimal. Only a very elementary part of the theory of weak convergence of probability measures will be needed; this is covered in the second part of Chapter 7. The techniques introduced here are of considerable importance beyond the needs of the book, since they are a foundation of the theory of approximation of random processes and limit theorems for sequences of random processes.

When one considers how stochastic approximation algorithms are used in applications, the fact of ultimate convergence with probability one can be misleading. Algorithms do not continue on to infinity, particularly when $\epsilon_n \rightarrow 0$. There is always a stopping rule that tells us when to stop the algorithm and to accept some function of the recent iterates as the “final value.” The stopping rule can take many forms, but whichever it takes, all that we know about the “final value” at the stopping time is information of a distributional type. There is no difference in the conclusions provided by the probability one and the weak convergence methods. In applications that are of concern over long time intervals, the actual physical model might “drift.” Indeed, it is often the case that the step size is not allowed to go to zero, and then there is no general alternative to the weak convergence methods at this time.

The ODE approach to the limit theorems obtains the ODE by appropriately averaging the dynamics, and then by showing that the limit set of the ODE is just the set of asymptotic points of the $\{\theta_n\}$. The ODE is easier to characterize, and requires weaker conditions and simpler proofs when weak convergence methods are used. Furthermore, it can be shown that $\{\theta_n\}$ spends “nearly all” of its time in an arbitrarily small neighborhood of the limit point or set. The use of weak convergence methods can lead to better probability one proofs in that, once we know that $\{\theta_n\}$ spends “nearly all” of its time (asymptotically) in some small neighborhood of the limit point, then a *local analysis* can be used to get convergence with probability one. For example, the methods of Chapters 5 and 6 can be applied locally, or the local large deviations methods of [39] can be used. Even when we can only prove weak convergence, if θ_n is close to a stable limit point at iterate n , then under broad conditions the mean escape time from a small neighborhood of that limit point is of the order of e^{c/ϵ_n} for some $c > 0$.

Section 7.2 is motivational in nature, aiming to relate some of the ideas of weak convergence to probability one convergence and convergence in distribution. It should be read “lightly.” The general theory is covered in Chapter 8 for a broad variety of algorithms, using what might be called “weak local ergodic theorems.” The essential conditions concern the rates of decrease of the conditional expectation of the future noise given the past noise, as the time difference increases. Chapter 9 illustrates the relative

convenience and power of the methods of Chapter 8 by providing proofs of convergence for some of the examples in Chapters 2 and 3.

Chapter 10 concerns the rate of convergence. Loosely speaking, a standard point of view is to show that a sequence of suitably normalized iterates, say, of the form $(\theta_n - \bar{\theta})/\sqrt{\epsilon_n}$ or $n^\beta(\theta_n - \bar{\theta})$ for an appropriate $\beta > 0$, converges in distribution to a random variable with mean zero and finite covariance matrix \bar{V} . We will do a little better and prove that the continuous time process obtained from suitably interpolated normalized iterates converges “weakly” to a stationary Gauss–Markov process, whose covariance matrix (at any time t) is \bar{V} . The methods use only the techniques of weak convergence theory, which are outlined in Chapter 7.

The use of stochastic approximation for the minimization of functions of a very high-dimensional argument has been of increasing interest. Owing to the high dimension, the classical Kiefer–Wolfowitz procedures can be very time consuming to use. As a result, there is much current interest in the so-called random-directions methods, where at each step n one chooses a direction d_n at random, obtains a noisy estimate \hat{Y}_n of the derivative in direction d_n , and moves an increment $-\epsilon_n \hat{Y}_n$. Although such methods have been of interest and used in various ways for a long time [99], convincing arguments concerning their value and the appropriate choices of the direction vectors and scaling were lacking. Recent work [166] proposed a different way of getting the directions and attracted needed attention to this problem. The proof of convergence of the random-directions methods that have been suggested to date are exactly the same as that for the classical Kiefer–Wolfowitz procedure (see Chapter 5). The comparison of the rates of convergence under the different ways of choosing the random directions is given at the end of Chapter 10, and shows that the older and newer methods have essentially the same properties, when the norms of the direction vectors d_n are the same. It is seen that the random-directions methods can be quite advantageous, but care needs to be exercised in their use.

The performance of the stochastic approximation algorithms depends heavily on the choice of the step size sequence ϵ_n , and the lack of a general approach to getting good sequences has been a handicap in applications. In [142], Polyak and Juditsky showed that, if the coefficients ϵ_n go to zero slower than $O(1/n)$, then the averaged sequence $\sum_{i=1}^n \theta_i/n$ converges to its limit at an optimal rate. This implies that the use of relatively large step sizes, while letting the “off-line” averaging take care of the increased noise effects, will yield a substantial overall improvement. These results have since been corroborated by numerous simulations and extended mathematically. In Chapter 11, it is first shown that the averaging improves the asymptotic properties whenever there is a “classical” rate of convergence theorem of the type derived in Chapter 10, including the constant $\epsilon_n = \epsilon$ case. This will give the minimal window over which the averaging will yield an improvement. The maximum window of averaging is then obtained by

a direct computation of the asymptotic covariance of the averaged process. Intuitive insight is provided by relating the behavior of the original and the averaged process to that of a three time scale discrete algorithm where it is seen that the key property is the separation of the time scales of the two slower processes.

Chapter 12 concerns the decentralized/asynchronous algorithms, where the work is split between several processors, each of which has control over a different set of parameters. The processors work at different speeds, and there can be delays in passing information to each other. Owing to the asynchronous property, the analysis must be in “real” rather than “iterate” time. This complicates the notation, but all of the results of the previous chapters can be carried over. Typical applications are decentralized optimization of queueing networks and Q -learning.

Some topics are not covered. As noted, the algorithm in continuous time differs little from that in discrete time. The basic idea can be extended to infinite-dimensional problems [10, 12, 41, 60, 104, 135, 148, 158, 161, 177, 178, 179, 202]. The function minimization problem where there are many local minima has attracted some attention [54, 96, 187], but little is known at this time concerning effective methods. Some effort has been devoted [18] to showing that suitable conditions on the noise guarantee that there cannot be convergence to an unstable or marginally stable point of the ODE. Such results are needed and do increase confidence in the algorithms. The conditions can be hard to verify, particularly in high-dimensional problems, and the results do not guarantee that the iterates would not actually spend a lot of time near such bad points, particularly when the step sizes are small and there is poor initial behavior. Additionally, one tries to design the procedure and to use variance reduction methods to reduce the effects of the noise. Perhaps the general problem of avoiding “bad” points needs to be looked at in the context of the “multiminimum” problem, or controlled perturbation methods used (say, analogous to that at the end of Chapter 5).

Penalty-multiplier and Lagrangian methods for constrained problems are discussed in [99]. They involve only minor variations on what is done here, but they are omitted for lack of space. We concentrate on algorithms defined on r -dimensional Euclidean space, except as modified by inequality or equality constraints. The treatment of the equality constrained problem shows that the theory also covers processes defined on smooth manifolds.

We express our deep gratitude to Paul Dupuis, Felisa Vazqu  z-Abad, and Sid Yakowitz for their careful reading and critical remarks on various parts of the manuscript. The long-term support and encouragement of the National Science Foundation and the Army Research Office are also gratefully acknowledged.

Notation and numbering. Chapters are divided into sections, and sections into subsections. Within a chapter, (1.2) (resp., (A2.1)) denotes Equa-

tion 2 of Section 1 (resp., Assumption 2 of Section 1). Section 1 (Subsection 1.2, resp.) always means the first section (resp., the second subsection of the first section) in the chapter in which the symbol is used. To refer to equations (resp., assumptions) in other chapters, we use, e.g., (1.2.3) (resp., (A1.2.3)) to denote the third equation (resp., the third assumption) in Section 2 of Chapter 1. When not in Chapter 1, Section 1.2 (resp., Subsection 1.2.3) means Section 2 (resp., Subsection 3 of Section 2) of Chapter 1.

Throughout the book, $\|\cdot\|$ denotes either a Euclidean norm or a norm on the appropriate function spaces, which will be clear from the context. A point x in a Euclidean space is a column vector, and the i th component of x is denoted by x_i . However, the i th component of θ is denoted by θ^i , since subscripts on θ are used to denote the value at a time n . The symbol $'$ denotes transpose. Moreover, both A' and $(A)'$ will be used interchangeably, e.g., both $g_n^{\epsilon, '}(\theta)$ and $(g_n^{\epsilon}(\theta))'$ denote the transpose of $g_n(\theta)$. Subscripts θ and x denote either a gradient or a derivative, depending on whether the variable is vector or real-valued. For convenience, we list many of the symbols at the end of the book.

Providence, Rhode Island
Detroit, Michigan

Harold J. Kushner
G. George Yin