Stochastic Approximation

This simple, compact toolkit for designing and analyzing stochastic approximation algorithms requires only basic literacy in probability and differential equations. Yet these algorithms have powerful applications – for example, in control and communications engineering, artificial intelligence and economic modelling.

The dynamical systems viewpoint treats an algorithm as a noisy discretization of a limiting differential equation and argues that, under reasonable hypotheses, it tracks the asymptotic behaviour of the differential equation with probability one. The limiting differential equation, which can usually be obtained by inspection, is easier to analyze.

Novel topics covered in the book include finite-time behaviour, multiple timescales and asynchronous implementation. A separate chapter gives a useful taxonomy of applications, with concrete examples from engineering and economics. Notably it covers several variants of stochastic gradient-based optimization schemes, fixed-point solvers, which are commonplace in learning algorithms for approximate dynamic programming, and some models of collective behaviour. Three appendices give self-contained summaries of background material from analysis, differential equations and probability.

Ideal for graduate students, researchers and practitioners in electrical engineering and computer science, especially those working in control, communications, signal processing and machine learning, it is also relevant to economics, probability and statistics.

Stochastic Approximation

A Dynamical Systems Viewpoint

Vivek S. Borkar

Tata Institute of Fundamental Research, Mumbai

CAMBRIDGE UNIVERSITY PRESS

HINDUSTAN BOOK AGENCY

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314-321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi - 110025, India

79 Anson Road, #06-04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9780521515924

Sold and distributed in all countries except India, Pakistan, Bangladesh and Sri Lanka by Cambridge University Press

Sold and distributed in India, Pakistan, Bangladesh and Sri Lanka by Hindustan Book Agency

First published 2008

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

ISBN 978-0-521-51592-4 Hardback

Cambridge University Press
978-0-521-51592-4 — Stochastic Approximation
Vivek S. Borkar
Frontmatter
More Information

Contents

Preface

Stochastic approximation was introduced in a 1951 article in the *Annals of Mathematical Statistics* by Robbins and Monro. Originally conceived as a tool for statistical computation, an area in which it retains a place of pride, it has come to thrive in a totally different discipline, viz., that of electrical engineering. The entire area of 'adaptive signal processing' in communication engineering has been dominated by stochastic approximation algorithms and variants, as is evident from even a cursory look at any standard text on the subject. Then there are the more recent applications to adaptive resource allocation problems in communication networks. In control engineering too, stochastic approximation is the main paradigm for on-line algorithms for system identification and adaptive control.

This is not accidental. The key word in most of these applications is *adaptive.* Stochastic approximation has several intrinsic traits that make it an attractive framework for adaptive schemes. It is designed for uncertain (read 'stochastic') environments, where it allows one to track the 'average' or 'typical' behaviour of such an environment. It is incremental, i.e., it makes small changes in each step, which ensures a graceful behaviour of the algorithm. This is a highly desirable feature of any adaptive scheme. Furthermore, it usually has low computational and memory requirements per iterate, another desirable feature of adaptive systems. Finally, it conforms to our anthropomorphic notion of adaptation: It makes small adjustments so as to improve a certain performance criterion based on feedbacks received from the environment.

For these very reasons, there has been a resurgence of interest in this class of algorithms in several new areas of engineering. One of these, viz., communication networks, is already mentioned above. Yet another major application domain has been artificial intelligence, where stochastic approximation has provided the basis for many learning or 'parameter tuning' algorithms in soft computing. Notable among these are the algorithms for training neural networks and the algorithms for reinforcement learning, a popular learning paradigm for autonomous software agents with applications in e-commerce, robotics, etc.

Yet another fertile terrain for stochastic approximation has been in the area of economic theory, for reasons not entirely dissimilar to those mentioned above. On one hand, they provide a good model for collective phenomena, where *micromotives* (to borrow a phrase from Thomas Schelling) of individual agents aggregate to produce interesting *macrobehaviour*. The 'nonlinear urn' scheme analyzed by Arthur and others to model increasing returns in economics is a case in point. On the other hand, their incrementality and low per iterate computational and memory requirements make them an ideal model of a *boundedly rational* economic agent, a theme which has dominated their application to learning models in economics, notably to learning in evolutionary games.

This flurry of activity, while expanding the application domain of stochastic approximation, has also thrown up interesting new issues, some of them dictated by technological imperatives. Consequently, it has spawned interesting new theoretical developments as well. The time thus seemed right for a book pulling together old and new developments in the subject with an eye on the aforementioned applications. There are, indeed, several excellent texts already in existence, many of which will be referenced later in this book. But they tend to be *comprehensive* texts: excellent for the already initiated but rather intimidating for someone who wants to make quick inroads. Hence a need for a 'bite-sized' text. The present book is an attempt at one.

Having decided to write a book, there was still a methodological choice. Stochastic approximation theory has two somewhat distinct strands of research. One, popular with statisticians, uses the techniques of martingale theory and associated convergence theorems for analysis. The second, popular more with engineers, treats the algorithm as a noisy discretization of an ordinary differential equation (*o.d.e.*) and analyzes it as such. We have opted for the latter approach, because the kind of intuition that it offers is an added advantage in many of the engineering applications.

Of course, this is not the first book expounding this approach. There are several predecessors such as the excellent texts by Benveniste–Metivier–Priouret, Duflo, and Kushner–Yin referenced later in the book. These are, however, what we have called *comprehensive* texts above, with a wealth of information. This book is *not* comprehensive, but is more of a compact account of the highlights to enable an interested, mathematically literate reader to run through the basic ideas and issues in a relatively short time span. The other 'novelties' of the book would be a certain streamlining and fine-tuning of proofs using that eternal source of wisdom – *hindsight*. There are occasional new variations on proofs sometimes leading to improved results (e.g., in Chapter 6) or just shorter proofs, inclusion of some newer themes in theory and applications, and so on. Given the nature of the subject, a certain mathematical sophistication was unavoidable. For the benefit of those not quite geared for it, we have collected the more advanced mathematical requirements in a few appendices. These should serve as a source for quick reference and pointers to the literature, but not as a replacement for a firm grounding in the respective areas. Such grounding is a must for anyone wishing to contribute to the *theory* of stochastic approximation. Those interested more in *applying* the results to their respective specialties may not feel the need to go much further than this little book.

Let us conclude this long preface with the pleasant task of acknowledging all the help received in this venture. The author forayed into stochastic approximation around 1993–1994, departing significantly from his dominant activity till then, which was controlled Markov processes. This move was helped by a project on adaptive systems supported by a Homi Bhabha Fellowship. More than the material help, the morale boost was a great help and he is immensely grateful for it. His own subsequent research in this area has been supported by grants from the Department of Science and Technology, Government of India, and was conducted in the two 'Tata' Institutes: Indian Institute of Science at Bangalore and the Tata Institute of Fundamental Research in Mumbai. Dr. V. V. Phansalkar went though the early drafts of a large part of the book and with his fine eye for detail, caught many errors. Prof. Shalabh Bhatnagar, Dr. Arzad Alam Kherani and Dr. Huizen (Janey) Yu also read the drafts and pointed out corrections and improvements (Janey shares with Dr. Phansalkar the rare trait for having a great eye for detail and contributed a lot to the final clean-up). Dr. Sameer Jalnapurkar did a major overhaul of chapters 1–3 and a part of chapter 4, which in addition to fixing errors, greatly contributed to their readability. Ms. Diana Gillooly of Cambridge University Press did an extremely meticulous job of editorial corrections on the final manuscript. The author takes full blame for whatever errors that remain. His wife Shubhangi and son Aseem have been extremely supportive as always. This book is dedicated to them.

Vivek S. Borkar
Mumbai, February 2008