Systems That Learn

An Introduction to Learning Theory for Cognitive and Computer Scientists

Daniel N. Osherson
Michael Stob
Scott Weinstein

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Times New Roman by Asco Trade Typesetting Ltd., Hong Kong, and printed and bound by Halliday Lithograph in the United States of America

Library of Congress Cataloging-in-Publication Data

Osherson, Daniel N.
Systems that learn.

(The MIT Press series in learning, development, and conceptual change)
"A Bradford book."
Bibliography: p.
Includes indexes.
1. Learning—Mathematical aspects. 2. Learning, Psychology of. 3. Human information processing—Mathematical models. I. Stob, Michael. II. Weinstein, Scott. III. Title. IV. Series.
BF318.083 1986 153.1'5'015113 85-19759
ISBN 0-262-15030-1

Systems That Learn

The MIT Press Series in Learning, Development, and Conceptual Change
Lila Gleitman, Susan Carey, Elissa Newport, and Elizabeth Spelke, editors

Names for Things: A Study in Human Learning, by John Macnamara, 1982

Conceptual Change in Childhood, by Susan Carey, 1985

"Gavagai!" or the Future History of the Animal Language Controversy, by David Premack, 1985

Systems That Learn: An Introduction to Learning Theory for Cognitive and Computer Scientists, by Daniel N. Osherson, Michael Stob, and Scott Weinstein, 1986

Series Foreword

This series in learning, development, and conceptual change will include state-of-the-art reference works, seminal book-length monographs, and texts on the development of concepts and mental structures. It will span learning in all domains of knowledge, from syntax to geometry to the social world, and will be concerned with all phases of development, from infancy through adulthood.

The series intends to engage such fundamental questions as

The nature and limits of learning and maturation: the influence of the environment, of initial structures, and of maturational changes in the nervous system on human development; learnability theory; the problem of induction; domain specific constraints on development.

The nature of conceptual change: conceptual organization and conceptual change in child development, in the acquisition of expertise, and in the history of science.

Lila Gleitman
Susan Carey
Elissa Newport
Elizabeth Spelke

Preface

It is a familiar observation that an organism's genotype may be conceived as a function that maps potential environments into potential phenotypes. Relativizing this conception to cognitive science allows human intellectual endowment to be construed as a particular function mapping early experience into mature cognitive competence. The function might be called "human nature relative to cognition." Learning theory is a mathematical tool for the study of this function. This book attempts to acquaint the reader with the use of this tool.

Less cryptically, learning theory is the study of systems that map evidence into hypotheses. Of special interest are the circumstances under which these hypotheses stabilize to an accurate representation of the environment from which the evidence is drawn. Such stability and accuracy are conceived as the hallmarks of learning. Within learning theory, the concepts "evidence," "stabilization," "accuracy," and so on, give way to precise definitions.

As developed in this book, learning theory is a collection of theorems about certain kinds of number-theoretic functions. We have discussed the application of such theorems to cognitive science and epistemology in a variety of places (e.g., Osherson, Stob, and Weinstein, 1984, 1985, 1985a; Osherson and Weinstein, 1982a, 1984, 1985). In contrast, the present work centers on the mathematical development of learning theory rather than on empirical hypotheses about human learning. As an aid to intuition, however, we have attempted to concretize the formal developments in this book through extended discussion of first language acquisition.

We have not tried to survey the immense field of machine inductive inference. Rather, we have selected for presentation just those results that seem to us to clarify questions relevant to human intellectual development. Several otherwise fascinating topics in machine learning have thus been left aside. Our choices no doubt reflect tacit theoretical commitments not universally shared. An excellent review of many topics passed over here is provided by Angluin and Smith (1982). Our own previously published work in the technical development of learning theory (e.g., Osherson and Weinstein, 1982, 1982a; Osherson, Stob, and Weinstein, 1982, 1982a, 1985) is entirely integrated herein.

Our concern in the present work for the mathematical development of learning theory has resulted in rigorous exposition. Less formal introductions to the central concepts and topics of learning theory are available in Osherson and Weinstein (1984) and Osherson, Stob, and Weinstein (1984).

We would be pleased to receive from our readers comments and corrections, as well as word of new results.

Acknowledgments

Our principal intellectual debts are to the works of E. Mark Gold and Noam Chomsky. Gold (1967) established the formal framework within which learning theory has developed. Chomsky's writings have revealed the intimate connection between the projection problem and human intelligence. In addition, we have been greatly influenced by the research of Blum and Blum (1975), Angluin (1980), Case and Smith (1983), and Wexler and Culicover (1980). Numerous conversations with Lila Gleitman and with Steven Pinker have helped us to appreciate the bearing of learning theory on empirical studies of first language acquisition, and conversely the bearing of first language acquisition studies on learning theory. We thank them for their patient explanations.

Preparation of the manuscript was facilitated by a grant to Osherson from the Fyssen Foundation for 1983–84 and by National Science Foundation Grants MCS 80-02937 and 82-00032 to Stob. We thank these agencies for their support.

How to Use This Book

Mathematical prerequisites for this text include elementary set theory and an intuitive understanding of the concept "computable function." Lewis and Papadimitriou (1981) provide an excellent introduction to this material. Acquaintance with the elementary portion of recursion theory is also advisable. We recommend Machtey and Young (1978).

Starred material in the text is of more advanced character and may be omitted without loss of continuity. We have relegated considerable exposition to the exercises, which should be at least attempted.

Definitions, examples, lemmas, propositions, open questions, and exercises are numbered independently within the section or subsection in which they appear. Thus proposition 4.4.1B refers to the second proposition of section 4.4.1; it appears before lemma 4.4.1A, the first lemma of the same section. Symbol, subject, and name indexes may be found at the end of the book.

We use standard set-theoretic notation and recursion-theoretic notation drawn from Rogers (1967) throughout. Note that $\subset$ denotes proper inclusion, whereas $\subseteq$ denotes (possibly improper) inclusion.

Systems That Learn

Contents

Let us play a game.

We have selected a set of numbers, and you must guess the set that we have in mind. The set consists of every positive integer with a sole exception. Thus the set might be $\{2, 3, 4, 5, \ldots\}$ or $\{1, 3, 4, 5, \ldots\}$ or $\{1, 2, 4, 5, \ldots\}$, etc. We will give you an unlimited number of clues about the set, and you are to guess after each clue. We will never tell you whether you are right.

First clue: The set contains the number 1.

Please guess the set we have in mind. Would you like to guess the set $\{2, 3, 4, 5, \ldots\}$? (That would be unwise.)

Second clue: The set contains the number 3.

Please make another guess. How about $\{1, 2, 3, 4, 5, 6, 8, 9, 10, \ldots\}$ or does that seem arbitrary to you?

Third clue: The set contains the number 4.

Go ahead and guess.

Fourth clue: The set contains the number 2.

Does the fourth clue surprise you? Guess again.

Fifth clue: The set contains the number 6.

Guess.

Sixth clue: The set contains the number 7.

Guess.

Seventh clue: The set contains the number 8.

Guess.

We interrupt the game at this point because we would like to ask you some questions about it.

First question: Are you confident about your seventh guess? Give an example of an eighth clue that would lead you to repeat your last guess. Give an example of an eighth clue that would lead you to change your guess.

Second question: Let us say that a "guessing rule" is a list of instructions for converting the clues received up to a given point into a guess about the set we have in mind. Were your guesses chosen according to some guessing rule, and if so, which one?

Third question: What should count as winning the game? Consider the following criterion: You win just in case at least one of your guesses is right. This criterion makes winning the game too easy. Say why.

Fourth question: We advocate the following criterion: You win just in case you eventually make the right guess and subsequently never change your mind regardless of the new clues you receive. In this case let us say that you "win in the limit." Is it possible to win the game in the limit even though you make one hundred wrong guesses? Is there any number of wrong guesses that is logically incompatible with winning the game in the limit?

Fifth question: Suppose that all the clues we give you are of the form: The set contains the number *n*. Suppose furthermore that for every positive integer *i*, we eventually give you a clue of this form if and only if *i* is in fact contained in the set we have in mind. (So for every number *i* in our set, you are eventually told that the set contains *i*; also you receive no false information about the set.) Do not suppose anything about the order in which you will get all these clues. We will order them any way we please. (Recall how we surprised you with the fourth clue.) Now let us call a guessing rule "winning" just in case the following is true. If you use the rule to choose your guesses, then no matter which of the sets we have in mind, you are guaranteed to win the game in the limit. Specify a winning guessing rule for our game.

Sixth question: We make the game harder. This time we are allowed to select any of the sets that are legal in the original game, but we may also select the set $\{1, 2, 3, 4, 5, 6, \dots\}$ of all positive integers. The rules about clues are the same as given in question 5. Play this new game with a friend, and then think about the following question. Is there a winning guessing rule for the new game?

Seventh question: Let us make the last game easier. The choice of sets is the same as in the last game, but we now agree to order our clues in a certain way. For all positive integers *i* and *j*, if both *i* and *j* are included in the set we have in mind, and if *i* is less than *j*, then you will receive the clue "The set contains *i*" before you receive the clue "The set contains *j*." Can you specify a winning guessing rule for this version of the game?

Eighth question: Here is another variant. We select a set from the original collection (thus the set $\{1, 2, 3, 4, 5, \dots\}$ of all positive integers is no longer allowed). Clues can be given in any order we please. You get only one guess. You may wait to see as many clues as you like, but your first guess is definitive. Play this game with a friend. Then show that no matter what rule you use to make your guess, you are not guaranteed to be right. Think about what happens if you are allowed two guesses in the game.

The games we have been playing resemble the process of scientific discovery. Nature plays our role, selecting a certain pattern that is imposed on the world. The scientist plays your role, examining an endless series