INTRODUCTION TO **DEEP LEARNING**

EUGENE CHARNIAK

Introduction to Deep Learning

Eugene Charniak

The MIT Press

This book was set in LaTeX by author. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data is available.
ISBN: 978-0-262-03951-2

10 9 8 7 6 5 4 3 2 1

To my family, once more

Contents

only in Chapter 1, when we build up to back-propagation from scratch and I would not be surprised if an extra lecture on partial derivatives would do. Last, there is a probability and statistics prerequisite. This simplifies the exposition and I certainly want to encourage students to take such a course. I also assume a rudimentary knowledge of programming in Python. I do not include this in the text, but my course has an extra "lab" on basic Python.

That your author was playing catch-up when writing this book also explains the fact that in almost every chapter's section on further reading you will find, beyond the usual references to important research papers, many reference to secondary sources — others' educational writings. I would never have learned this material without them.

Providence, Rhode Island
January 2018