

DEEP LEARNING

JOHN D. KELLEHER

The MIT Press Essential Knowledge Series

A complete list of the titles in this series appears at the back of this book.

The MIT Press | Cambridge, Massachusetts | London, England

CONTENTS

Series Foreword	vii
Preface	ix
Acknowledgments	xi

© 2019 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Chaparral Pro by Toppan Best-set Premedia Limited.
Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Kelleher, John D., 1974- author.

Title: Deep learning / John D. Kelleher.

Description: Cambridge, MA : The MIT Press, [2019] | Series:

The MIT press essential knowledge series | Includes bibliographical references and index.

Identifiers: LCCN 2018059550 | ISBN 9780262537551 (pbk. : alk. paper)

Subjects: LCSH: Machine learning. | Artificial intelligence.

Classification: LCC Q325.5 .K454 2019 | DDC 006.3/1—dc23 LC record available at <https://lcn.loc.gov/2018059550>

1098765

1	Introduction to Deep Learning	1
2	Conceptual Foundations	39
3	Neural Networks: The Building Blocks of Deep Learning	65
4	A Brief History of Deep Learning	101
5	Convolutional and Recurrent Neural Networks	159
6	Learning Functions	185
7	The Future of Deep Learning	231

Glossary 251

Notes 257

References 261

Further Readings 267

Index 269

NEURAL NETWORKS: THE BUILDING BLOCKS OF DEEP LEARNING

The term *deep learning* describes a family of *neural network* models that have multiple layers of simple information processing programs, known as neurons, in the network. The focus of this chapter is to provide a clear and comprehensive introduction to how these neurons work and are interconnected in artificial neural networks. In later chapters, we will explain how neural networks are trained using data.

A neural network is a computational model that is inspired by the structure of the human brain. The human brain is composed of a massive number of nerve cells, called neurons. In fact, some estimates put the number of neurons in the human brain at one hundred billion (Herculano-Houzel 2009). Neurons have a simple three-part structure consisting of: a cell body, a set of fibers called dendrites, and a single long fiber called an axon.

Figure 3.1 illustrates the structure of a neuron and how it connects to other neurons in the brain. The dendrites and the axon stem from the cell body, and the dendrites of one neuron are connected to the axons of other neurons. The dendrites act as input channels to the neuron and receive signals sent from other neurons along their axons. The axon acts as the output channel of a neuron, and so other neurons, whose dendrites are connected to the axon, receive the signals sent along the axon as inputs.

Neurons work in a very simple manner. If the incoming stimuli are strong enough, the neuron transmits an electrical pulse, called an action potential, along its axon to the other neurons that are connected to it. So, a neuron acts as an all-or-none switch, that takes in a set of inputs and either outputs an action potential or no output.

This explanation of the human brain is a significant simplification of the biological reality, but it does capture




Figure 3.1 The structure of a neuron in the brain.

the main points necessary to understand the analogy between the structure of the human brain and computational models called neural networks. These points of analogy are: (1) the brain is composed of a large number of interconnected and simple units called neurons; (2) the functioning of the brain can be understood as processing information, encoded as high or low electrical signals, or activation potentials, that spread across the network of neurons; and (3) each neuron receives a set of stimuli from its neighbors and maps these inputs to either a high- or low-value output. All computational models of neural networks have these characteristics.

Artificial Neural Networks

An artificial neural network consists of a network of simple information processing units, called neurons. The power of neural networks to model complex relationships is not the result of complex mathematical models, but rather emerges from the interactions between a large set of simple neurons.

Figure 3.2 illustrates the structure of a neural network. It is standard to think of the neurons in a neural network as organized into layers. The depicted network has five layers: one input layer, three hidden layers, and one output layer. A hidden layer is just a layer that is neither