

Deep Learning

A Practical Introduction

Manel Martínez-Ramón
Meenu Ajith
Aswathy Rajendra Kurup



WILEY

Deep Learning

Deep Learning

A Practical Introduction

Manel Martínez-Ramón

Department of Electrical and Computer Engineering, The University of New Mexico,
Albuquerque, NM, USA

Meenu Ajith

Tri-Institutional Center for Translational Research in Neuroimaging and Data Science
(TReNDS), Georgia State University, Georgia Institute of Technology, and Emory
University, Atlanta, GA, USA

Aswathy Rajendra Kurup

Machine learning Engineer, Intel Corporation, Hillsboro, OR, USA

WILEY

This edition first published 2024.

© 2024 John Wiley & Sons Ltd

All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by law. Advice on how to obtain permission to reuse material from this title is available at <http://www.wiley.com/go/permissions>.

The right of Manel Martínez-Ramón, Meenu Ajith, and Aswathy Rajendra Kurup to be identified as the authors of this work has been asserted in accordance with law.

Registered Offices

John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, USA

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, customer services, and more information about Wiley products visit us at www.wiley.com.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty

While the publisher and authors have used their best efforts in preparing this work, they make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Library of Congress Cataloging-in-Publication Data applied for:

Hardback ISBN: 9781119861867

Cover Design: Wiley

Cover Image: © Yuichiro Chino/Getty Images

Set in 9.5/12.5pt STIXTwoText by Straive, Chennai, India

To our families, who have been unwavering in their support and understanding throughout the long nights and weekends spent on this journey. Your love and encouragement have fueled our passion for the world of deep learning, and this book is dedicated to you with profound gratitude.

To all the dreamers, may this book inspire you to chase your passions and never give up. Keep reaching for the stars.

To our shared dreams, relentless passion, and enduring friendship – this book is a testament to our collective journey.

Contents

	About the Authors	<i>xv</i>
	Foreword	<i>xvii</i>
	Preface	<i>xix</i>
	Acknowledgment	<i>xxi</i>
	About the Companion Website	<i>xxiii</i>
1	The Multilayer Perceptron	1
1.1	Introduction	1
1.2	The Concept of Neuron	2
1.2.1	The Perceptron	4
1.2.2	The Perceptron (Training) Rule	6
1.2.3	The Minimum Mean Square Error Training Criterion	8
1.2.4	The Least Mean Squares Algorithm	13
1.3	Structure of a Neural Network	14
1.3.1	The Multilayer Perceptron	17
1.3.2	Multidimensional Array Multiplications	19
1.4	Activations	21
1.5	Training a Multilayer Perceptron	22
1.5.1	Maximum Likelihood Criterion	22
1.5.2	Activations and Likelihood Functions	24
1.5.2.1	Logistic Activation for Binary Classification	24
1.5.2.2	Softmax Activation for Multiclass Classification	26
1.5.2.3	Gaussian Activation in Regression	28
1.5.3	The Backpropagation Algorithm	29
1.5.3.1	Gradient with Respect to the Output Weights	29
1.5.3.2	Gradient with Respect to Hidden Layer Weights	31
1.5.4	Summary of the BP Algorithm	34
1.6	Conclusion	37
	Problems	37
2	Training Practicalities	41
2.1	Introduction	41
2.2	Generalization and Overfitting	42

2.2.1	Basic Weight Initializations	43
2.2.2	Activation Aware Initializations	44
2.2.3	MiniBatch Gradient Descent	44
2.3	Regularization Techniques	45
2.3.1	L_1 and L_2 Regularization	46
2.3.2	Dropout	47
2.3.3	Early Stopping	48
2.3.4	Data Augmentation	48
2.4	Normalization Techniques	50
2.5	Optimizers	52
2.5.1	Momentum Optimization	53
2.5.2	Nesterov-Accelerated Gradient	54
2.5.3	AdaGrad	54
2.5.4	RMSProp	55
2.5.5	Adam	55
2.5.6	Adamax	56
2.6	Conclusion	58
	Problems	59
3	Deep Learning Tools	61
3.1	Python: An Overview	61
3.1.1	Variables	62
3.1.2	Statements, Indentation, and Comments	65
3.1.3	Conditional Statements	66
3.1.4	Loops	67
3.1.5	Functions	69
3.1.6	Objects and Classes	69
3.2	NumPy	72
3.2.1	Installation and Importing NumPy Package	72
3.2.2	NumPy Array	72
3.2.3	Creating Different Types of Arrays	74
3.2.4	Manipulating Array Shape	75
3.2.5	Stacking and Splitting NumPy Arrays	76
3.2.6	Indexing and Slicing	78
3.2.7	Arithmetic Operations and Mathematical Functions	79
3.3	Matplotlib	83
3.3.1	Plotting	83
3.3.1.1	Functional Method	83
3.3.1.2	Object Oriented Method	84
3.3.2	Customized Plotting	85
3.3.3	Two-dimensional Plotting	86
3.3.3.1	Bar Plot	87
3.3.3.2	Histogram	88
3.3.3.3	Pie Plot	89
3.3.3.4	Scatter Plot	89

- 3.3.3.5 Quiver Plot 90
- 3.3.3.6 Contour Plot 91
- 3.3.3.7 Box Plot 91
- 3.3.3.8 Violin Plot 92
- 3.3.4 Three-dimensional Plotting 93
 - 3.3.4.1 3D Contour 93
 - 3.3.4.2 3D Surface 94
 - 3.3.4.3 3D Wireframe 95
- 3.4 Scipy 97
 - 3.4.1 Data Input–Output Using Scipy 97
 - 3.4.2 Clustering Methods 98
 - 3.4.3 Constants 99
 - 3.4.4 Linear Algebra and Integration Routines 99
 - 3.4.5 Optimization 101
 - 3.4.6 Interpolation 102
 - 3.4.7 Image Processing 105
 - 3.4.8 Special Functions 106
- 3.5 Scikit-Learn 107
 - 3.5.1 Scikit-Learn API 107
 - 3.5.1.1 Estimator Interface 107
 - 3.5.1.2 Predictor Interface 107
 - 3.5.1.3 Transformer Interface 107
 - 3.5.2 Loading Datasets 108
 - 3.5.3 Data Preprocessing 109
 - 3.5.4 Feature Selection 113
 - 3.5.5 Supervised and Unsupervised Learning Models 114
 - 3.5.6 Model Selection and Evaluation 115
- 3.6 Pandas 116
 - 3.6.1 Pandas Data Structures 117
 - 3.6.1.1 Series 117
 - 3.6.1.2 Dataframe 117
 - 3.6.2 Data Selection 118
 - 3.6.3 Data Manipulation 118
 - 3.6.3.1 Sorting 118
 - 3.6.3.2 Grouping 119
 - 3.6.4 Handling Missing Data 120
 - 3.6.5 Input–Output Tools 121
 - 3.6.6 Data Information Retrieval 122
 - 3.6.7 Data Operations 122
 - 3.6.8 Data Visualization 123
- 3.7 Seaborn 125
 - 3.7.1 Seaborn Datasets 125
 - 3.7.2 Plotting with Seaborn 126
 - 3.7.2.1 Univariate Plots 126
 - 3.7.2.2 Bivariate Plots 126

3.7.2.3	Multivariate Plots	127
3.7.3	Additional Plotting Functions	129
3.7.3.1	Correlation Plots	129
3.7.3.2	Point Plots	130
3.7.3.3	Cat Plots	130
3.8	Python Libraries for NLP	131
3.8.1	Natural Language Toolkit (NLTK)	131
3.8.2	SpaCy	132
3.8.3	NLP Techniques	132
3.8.3.1	Tokenization	133
3.8.3.2	Stemming	135
3.8.3.3	Lemmatization	136
3.8.3.4	Stop Words	137
3.9	TensorFlow	138
3.9.1	Introduction	138
3.9.2	Elements of Tensorflow	139
3.9.3	TensorFlow Pipeline	139
3.10	Keras	141
3.10.1	Introduction	141
3.10.2	Elements of Keras	141
3.10.2.1	Models	142
3.10.2.2	Layers	142
3.10.2.3	Core Modules	142
3.10.3	Keras Workflow	142
3.11	Pytorch	144
3.11.1	Introduction	144
3.11.2	Elements of PyTorch	145
3.11.2.1	PyTorch Tensors	145
3.11.2.2	PyTorch Variables	146
3.11.2.3	Dynamic Computational Graphs	146
3.11.2.4	Modules	146
3.11.3	Workflow of Pytorch	147
3.12	Conclusion	149
	Problems	150
4	Convolutional Neural Networks	153
4.1	Introduction	153
4.2	Elements of a Convolutional Neural Network	153
4.2.1	Overall Structure of a CNN	154
4.2.2	Convolutions	155
4.2.3	Convolutions in Two Dimensions	156
4.2.4	Padding	158
4.2.5	Stride	159
4.2.6	Pooling	160
4.3	Training a CNN	160

4.3.1	Formulation of the Convolution Layer in a CNN	160
4.3.2	Backpropagation of a Convolution Layer	162
4.3.3	Forward Step in a CNN	163
4.3.4	Backpropagation in the Dense Section of a CNN	164
4.3.5	Backpropagation of the Convolutional Section of a CNN	164
4.4	Extensions of the CNN	166
4.4.1	AlexNet	166
4.4.2	VGG	168
4.4.3	Inception	169
4.4.4	ResNet	170
4.4.5	Xception	171
4.4.6	MobileNet	172
4.4.6.1	Depthwise Separable Convolutions	173
4.4.6.2	Width Multiplier	174
4.4.6.3	Resolution Multiplier	174
4.4.7	DenseNet	174
4.4.8	EfficientNet	176
4.4.9	Transfer Learning for CNN Extensions	177
4.4.10	Comparisons Among CNN Extensions	181
4.5	Conclusion	184
	Problems	184
5	Recurrent Neural Networks	187
5.1	Introduction	187
5.2	RNN Architecture	188
5.2.1	Structure of the Basic RNN	188
5.2.2	Input–Output Configurations	190
5.3	Training an RNN	191
5.3.1	Gradient with Respect to the Output Weights	194
5.3.2	Gradient with Respect to the Input Weights	195
5.3.3	Gradient with Respect to the Hidden State Weights	196
5.3.4	Summary of the Backpropagation Through Time in an RNN	196
5.4	Long-Term Dependencies: Vanishing and Exploding Gradients	199
5.5	Deep RNN	201
5.6	Bidirectional RNN	203
5.7	Long Short-Term Memory Networks	204
5.7.1	LSTM Gates	205
5.7.2	LSTM Internal State	205
5.7.3	Hidden State and Output of the LSTM	206
5.7.4	LSTM Backpropagation	208
5.7.5	Machine Translation with LSTM	210
5.7.6	Beam Search in Sequence to Sequence Translation	212
5.8	Gated Recurrent Units	218
5.9	Conclusion	221
	Problems	222

6	Attention Networks and Transformers	225
6.1	Introduction	225
6.2	Attention Mechanisms	227
6.2.1	The Nadaraya–Watson Attention Mechanism	227
6.2.2	The Bahdanau Attention Mechanism	229
6.2.3	Attention Pooling	232
6.2.4	Representation by Self-Attention	233
6.2.5	Training the Self-Attention Parameters	234
6.2.6	Multi-head Attention	235
6.2.7	Positional Encoding	236
6.3	Transformers	242
6.4	BERT	249
6.4.1	BERT Architecture	250
6.4.2	BERT Pre-training	250
6.4.3	BERT Fine-Tuning	252
6.4.4	BERT for Different NLP Tasks	252
6.5	GPT-2	256
6.5.1	Language Modeling	257
6.6	Vision Transformers	262
6.6.1	Comparison between ViTs and CNNs	264
6.7	Conclusion	269
	Problems	270
7	Deep Unsupervised Learning I	273
7.1	Introduction	273
7.2	Restricted Boltzmann Machines	274
7.2.1	Boltzmann Machines	274
7.2.2	Training a Boltzmann Machine	275
7.2.3	The Restricted Boltzmann Machine	276
7.3	Deep Belief Networks	278
7.3.1	Training a DBN	278
7.4	Autoencoders	279
7.4.1	Autoencoder Framework	279
7.5	Undercomplete Autoencoder	284
7.6	Sparse Autoencoder	285
7.7	Denoising Autoencoders	287
7.7.1	Denoising Autoencoder Algorithm	287
7.8	Convolutional Autoencoder	288
7.9	Variational Autoencoders	291
7.9.1	Latent Variable Inference: Lower Bound Estimation Approach	292
7.9.2	Reparameterization Trick	294
7.9.3	Illustration: Variational Autoencoder Implementation	295
7.10	Conclusion	297
	Problems	298

8	Deep Unsupervised Learning II	301
8.1	Introduction	301
8.2	Elements of GAN	303
8.2.1	Generator	304
8.2.2	Discriminator	304
8.3	Training a GAN	305
8.4	Wasserstein GAN	309
8.5	DCGAN	312
8.5.1	DCGAN Training and Outcomes Highlights	313
8.6	cGAN	316
8.6.1	cGAN Training and Outcomes Highlights	318
8.7	CycleGAN	318
8.7.1	CycleGAN Training and Outcomes Highlights	321
8.7.2	Applications of CycleGAN	323
8.8	StyleGAN	323
8.8.1	StyleGAN Properties and Outcome Highlights	326
8.9	StackGAN	328
8.9.1	StackGAN Training and Outcomes Highlights	331
8.10	Diffusion Models	333
8.10.1	Forward Diffusion Process	334
8.10.2	Reverse Diffusion Process	335
8.10.3	Diffusion Process Training	335
8.11	Conclusion	338
	Problems	339
9	Deep Bayesian Networks	341
9.1	Introduction	341
9.2	Bayesian Models	342
9.2.1	The Bayes' Rule	342
9.2.2	Priors as Regularization Criteria	343
9.3	Bayesian Inference Methods for Deep Learning	344
9.3.1	Markov Chain Monte Carlo Methods	344
9.3.2	Hamiltonian MCMC	347
9.3.3	Variational Inference	349
9.3.4	Bayes by Backpropagation	351
9.4	Conclusion	352
	Problems	353
	List of Acronyms	355
	Notation	359
	Bibliography	365
	Index	387

About the Authors

Dr. Manel Martínez-Ramón received his Telecommunication Engineering degree from Universitat Politècnica de Catalunya, Spain, in 1994 and his PhD in Telecommunication Engineering from Universidad Carlos III de Madrid, Spain, in 1999. He is currently a professor of Artificial Intelligence with the Department of Electrical and Computer Engineering of the University of New Mexico, NM, USA, where he holds the King Felipe VI Endowed Chair. His research interests are in the area of machine learning, where he has produced numerous contributions to kernel learning methods, Gaussian processes, and deep learning, with applications to electromagnetics and antenna array processing, smart grid, scientific particle accelerators, and others. As an instructor, he teaches graduate courses in statistical learning theory, Gaussian process learning, probabilistic machine learning, and deep learning both face-to-face and online.

Dr. Meenu Ajith earned her PhD in Electrical Engineering from the University of New Mexico, USA, in 2022. Presently, she serves as a postdoctoral research associate at the Tri-Institutional Center for Translational Research in Neuroimaging and Data Science (TReNDS), a collaborative research institute supported by Georgia State, the Georgia Institute of Technology, and Emory University in Atlanta, GA, USA. Her research focuses on deep learning, image processing, time-series analysis, and neuroimaging. She obtained her MS in Electrical Engineering from the University of New Mexico in 2017 and her bachelor's degree in Electronics and Communication Engineering from Amrita School of Engineering in 2015. In her current role as a postdoctoral researcher, Dr. Ajith concentrates on implementing and applying various deep-learning models and neuroinformatic tools that leverage advanced brain imaging data. Her objective is to translate these approaches into biomarkers, addressing pertinent aspects of brain health and diseases.

Dr. Aswathy Rajendra Kurup earned her PhD in Electrical Engineering from the University of New Mexico USA in the year 2022, where her research focused on designing CNN-based deep-learning models for applications such as smart grids, medical diagnosis, and computer vision. She completed her MS in Electrical Engineering from the University of New Mexico, USA, in 2017. Currently, she serves as a Machine Learning Engineer at Intel Corporation, a leading force in the semiconductor chip manufacturing landscape.

In her current role, she applies her extensive knowledge to address real-world challenges, utilizing her expertise in handling diverse data types, including images, videos, and time-series data. As a part of the role, she applies data mining and statistical modeling techniques along with developing Machine learning/Deep learning solutions for enabling factory decision-making, improved equipment performance, and higher product yields.

Foreword

Deep Learning: A Practical Introduction, authored by Manel Martínez-Ramón, Meenu Ajith, and Aswathy Rajendra Kurup, stands as a pragmatic guide, which prepares the engaged student to digest and understand advanced deep learning concepts. Designed primarily as an educational resource for graduate-level courses in deep learning, this book is enriched with a valuable collection of exercises and practical Python tutorials, making it an ideal educational tool.

Deep learning, a cornerstone of modern artificial intelligence, has seen a meteoric rise in usage, powering the creation of text, images, and videos, from simple prompts, and enhancing our predictive capabilities in a diverse array of applications. This book offers a thorough exploration of deep learning fundamentals, an essential component for students in engineering or computer science.

The authors begin by tracing the intriguing history of deep learning, setting the stage for a deeper dive into the subject. They skillfully introduce various methods for training and optimizing algorithms, alongside an overview of essential programming tools and libraries which are prevalent today, including Python, NumPy, TensorFlow, and Pytorch.

The book then covers a broad range of fundamental models including recurrent neural networks, transformers, unsupervised learning, and deep Bayesian networks. Within each of these chapters, there is an accessible introduction and detailed explanation of each modeling framework, which allows the reader who is new to deep learning to gain a foothold in this extraordinarily important space, while also providing practical examples including code and data as well as references for further learning. Additionally, it offers references for extended learning, bridging the gap between fundamental concepts and recent advancements in the field.

The author's provides a clear and comprehensive introduction to deep learning, making it an essential addition to the field's literature. Whether you are an instructor designing a course or a student embarking on self-directed learning, this book is an invaluable resource for navigating the complexities and applications of deep learning.

In essence, *Deep Learning: A Practical Introduction* is not just a textbook; it is a gateway to understanding and applying one of the most influential technologies in the field of artificial intelligence today. It is a useful tool for (i) instructors who want to teach core deep learning topics to their students, (ii) researchers in a variety of fields, including my own field of neuroimaging, who want to develop domain-specific methods, and (iii) students who are interested in self-learning on this important topic.

Overall, I strongly endorse *Deep Learning: A Practical Introduction* as a valuable resource for both educators aiming to impart core deep learning concepts to their students and for learners pursuing self-study in this vital area. The book's blend of theoretical insights and practical applications, including code and data examples, makes it a standout choice for anyone looking to delve into the world of deep learning.

Vince Calhoun

Preface

The present book is intended to be a comprehensive introduction to deep learning that covers all major areas in this discipline. This document is designed to cover a full semester graduate class in deep learning, and it contains all the materials necessary to build the class. We structured our work in a classical way, starting from the fundamentals of neural networks, which are then used to describe the different elements of deep learning used in artificial intelligence, from the classic convolutional neural network and recurrent neural networks (RNNs) to the transformers, plus unsupervised learning structures and algorithms. In every chapter, we follow a schema where first the structures are described, and then the criteria and algorithms to optimize them are developed. In most cases, full mathematical developments are included in the description of the structure optimization.

Chapter 1 is a first contact with deep learning, where we introduce the most basic type of feedforward neural network (FFNN), which is called the multilayer perceptron (MLP). Here, we first introduce the low-level basic elements of most neural networks and then the structure and learning criteria.

Chapter 2 is complementary to Chapter 1, but its contents are valid for the rest of the book since it provides details about the practical training of deep learning structures, which we have omitted from the first chapter in order to make it more concise and compact.

These readers who do not have a knowledge of basic Python will benefit from using Chapter 3 in order to start experimenting with learning machines in this programming language. In this chapter, authors assume that the reader has reviewed Chapter 1, which implies that they have been introduced to the concepts of structure, criteria, and algorithms. If so, readers already had the opportunity to see some basic Python codes containing at least a class with methods and an instantiation of it to be used in the examples and exercises, without needing to understand their Python structure. In this chapter, we introduce the basic elements of Python to be used throughout the book, and we will revisit the code previously introduced in Chapter 3, among other examples.

The concepts and structure of convolutional neural structures are described in Chapter 4. It starts with the concept of convolution in two dimensions and its justification for its use in deep learning, after which the structure of a convolutional neural network is described. The training of such a structure is not commonly found in the literature, assuming that the students and practitioners understand and can apply the backpropagation to them. We offer in this chapter a full development of the backpropagation for convolutional neural networks and we summarize the algorithms, so the practitioner can program it. Still, most importantly, they will understand exactly how it works.

Chapter 5 covers the basics of the RNN. The chapter starts off with the architecture of the RNN and then explains how these networks are used for modeling sequential information. Further into the chapter, the training criterion is introduced, which describes the feed-forward training, loss functions, and backpropagation through time. Next, the different types of RNN and their application are discussed. The following section explains the shortcomings of RNNs and highlights the details on different types of gradient problems and the solutions to these problems. Then, the shortcomings of RNNs and highlights the details on different types of gradient problems and the solutions to these problems are explained. After that, the details on other RNN-derived structures which were introduced to mitigate the short-term memory problem associated with the traditional RNNs are discussed.

Chapter 6 provides a structured and comprehensive overview of the developments in attention-based networks. The first section summarizes the different types of attention mechanisms based on sequence, levels, positions, and representations. Finally, we review the network architectures that widely use attention and also discuss a few applications in which attention-based networks have shown a significant impact.

Chapter 7 gives a comprehensive outline of deep unsupervised learning. The overview gives an introduction to the two main categories of deep unsupervised learning such as probabilistic and nonprobabilistic models. The chapter is mainly devoted to the autoencoder, which is one of the widely used nonprobabilistic deep unsupervised learning methods. First, the basic elements, training criteria, and the extensions of autoencoders are explained. Following this, an overview of the deep belief networks (DBNs) is given and it constitutes the basic blocks (restricted Boltzmann machines), training using contrastive divergence, and the variations of DBN. Finally, we also provide different applications of unsupervised deep learning.

Chapter 8 briefly covers the generative adversarial networks (GANs). Primarily, it introduces the two elements of GANs namely discriminator and generator. After this, the complete architecture of the GAN is illustrated to have a higher level of understanding of the network. Next, the training criteria are outlined which describes the alternate training process between the discriminator and the generator. The loss functions that model the probability distribution of the data is also added in this section. Finally, popular models derived from GAN are presented, and the chapter is concluded by summarizing the advantages and trade-offs of GAN.

Chapter 9 covers the main topics of deep Bayesian networks. Here, the authors do not intend to be exhaustive by covering the state of the art of deep Bayesian networks, Instead, we propose a chapter that gives the reader a general view of the characteristics and different philosophies of Bayesian networks with respect to previously introduced structures and algorithms. After introducing the general concepts of deep Bayesian networks, including structures and criteria (thus following the same format used in the rest of the book) we explain the main optimization algorithms used in the current literature, with several examples.

June, 2024
Albuquerque, New Mexico

Manel Martínez-Ramón
Meenu Ajith
Aswathy Rajendra Kurup

Acknowledgment

Manel Martínez-Ramón has been partially supported by the King Felipe VI Endowed Chair of the University of New Mexico, NM, USA.

About the Companion Website

A repository in GitHub with the URL

<https://github.com/DeepLearning-book>



contains all the additional materials of this book. In particular, readers will find:

- The Python code (in Jupyter Notebook format) of all the examples provided throughout the book, so that the student or the practitioner can run them immediately.
- A complete set of slides written in LaTeX that summarize all chapters, intended to help instructors in the development of their lectures. The source files are also available so that instructors can modify the material and adapt it to each particular course design. All materials are available in the repository.

1

The Multilayer Perceptron

1.1 Introduction

The concept of artificial intelligence (AI) is relatively simple to explain, and it can be enunciated as a possible answer to the question of how to make a machine that is able to perform a given task without being explicitly programmed for it, but instead, extracting the necessary information from a set of data. Let us say, for example, that a machine is needed to classify green and red apples. The machine is provided with a camera, and all the mechanisms necessary to place one apple at a time in front of it and then throw it in one of two buckets. A machine wired to do this will relay in binary operators as “IF,” “THEN.” If the color is red, throw it in bucket A, otherwise, in bucket “B.”

The limitations of this method are obvious. If a pear is mistakenly introduced in the process, it will be classified as a green apple. Also, how can we use the same or similar structure for a different or more complex task? As in the previous machine, an AI approach uses features found in the data in order to take the decision, but the algorithm is not explicitly programmed. Instead, the machine has a specific parametric structure capable of learning from data. The learning process involves the optimization of a certain measurable criterion with respect to the parameters. The deep learning (DL) structures for artificial intelligence are able to learn complex tasks from the available data, but they also have capabilities such as learning how to extract the useful features for the task at hand, provide probabilistic outputs (i.e. “the probability of apple is 97%”), and many others. The basic element of such a structure in DL is the so-called artificial neuron, a simple concept that provides the power and nonlinear properties.

This chapter is intended to be a first contact with DL, where we introduce the most basic type of feedforward neural network (FFNN), which is called the multilayer perceptron (MLP). Here, we first introduce the low-level basic elements of most neural network (NN)s, then the structure and learning criteria.

The elements introduced in this chapter will be used throughout the book. We start from the single perceptron, we construct a basic MLP, where the different activations are developed, and then the notation based on tensors is also justified as a generalized tool to be used throughout the book. After this, we present the maximum likelihood (ML) criterion as a general criterion, which is then particularized to the classic cases corresponding to the

different output activations. Finally, the backpropagation (BP) is detailed and then summarized so that can be translated into a computer program.

In this chapter, examples and exercises are presented in a way that assumes that the student does not necessarily know about programming in Python. Examples will be focused on the behavior of the MLP, without focusing on the programming, and the exercises intended to modify, at a high-level data, parameters, and structures in order to answer questions to different practical cases. Chapter 3 explains, in particular, how the different examples have been coded, thus they will be reviewed in that chapter from the point of view of practical programming.

1.2 The Concept of Neuron

The idea of the artificial neural network (ANN) is obviously inspired by the structure of the nervous system. The first attempt to understand how neural tissue works from a logical perspective was published in 1943 by Warren S. McCulloch and Walter Pitts (1943) (Fig. 1.1). They proposed the first mathematical model for a biological neuron in his paper. In this model, the neuron has two possible states, defined as 0 or 1 depending on whether the neuron is resting or it has been activated or *fired*. This represents the *axon* of the neuron. The input of this neuron model consists of a number of *dendrites* whose excitation is also binary. This elemental structure is completed with an inhibitory input. If this input




Figure 1.1 Warren S. McCulloch (left) and Walter Pitts in 1949. Source: R. Moreno-Díaz and A. Moreno-Díaz (2007)/with permission from Elsevier.

is activated, the neuron cannot fire. If the inhibitory input is deactivated, the neuron can be activated if the combination of inputs is larger than a given threshold. This model is fully binary and, since it includes mathematical functions that cannot be differentiated, it cannot be treated mathematically in an easy way. Certain modifications that will be seen further give rise to what is known as the artificial neuron in use today.

Section 1.2.1 contains an introduction to the concept of artificial perceptron from an algebraic point of view. A possible way to train a single perceptron is introduced in Sections 1.2.2 and 1.2.3, as well as the limitations of this structure as a linear classifier.

The concept of artificial NN was introduced by the psychologist Frank Rosenblatt (Fig. 1.2) in 1958 Rosenblatt (1957, 1958). In this paper, he proposed a structure of the visual cortex *perceptron* (Fig. 1.3). The structure presented in Rosenblatt (1958) contained the fundamental idea that is used in any artificial learning structure. In the first stage




Figure 1.2 Frank Rosenblatt. Source: <https://news.cornell.edu/stories/2019/09/professors-perceptron-paved-way-ai-60-years-too-soon/> last accessed November 30, 2023.




Figure 1.3 The perceptron as described in Rosenblatt (1958)/American Psychological Association.

(Retina), the device collects the available observation or input pattern intended to be processed in order to extract knowledge of it. The second stage (Projection area) is in charge of processing this observation to extract the information needed for the task at hand. This information is commonly called the set of features of the input pattern. The third stage (Association area) is intended to process these features to map them into a given response. For example, the response may be to recognize some given object classes present in the scene. Rosenblatt is the father of the artificial perceptron. He proved that by modifying the McCulloch–Pitts neuron model, the neuron could actually learn tasks from the data. In particular, his model had weights that multiplied each of the inputs to the neuron as well as the input bias or threshold that could be adjusted for the neuron to perform a given task. He developed the Mark 1 perceptron machine, which was the first implementation of his perceptron algorithm. This device was not a computer but an electromechanical *learning machine*. The machine consisted of a camera constructed with an array of 400 photocells, the output of each one connected randomly to the dendrites of a set of neurons. The weights, or attenuations applied to these inputs, were controlled with potentiometers whose axes were connected to electric motors. During the learning procedure, the motors adjusted the input weights. This machine was able to distinguish *linearly separable* patterns, or patterns that were at one or another side of a hyperplane in the space of 400 dimensions spanned by the camera inputs depending on its binary class. The invention was then limited in its capabilities until it was proven that a perceptron constructed with more than one layer of neurons MLP had nonlinear capabilities, that is, the ability to separate patterns that could not be separated by a hyperplane. Nevertheless, the MLP could not be trained using the techniques introduced by Rosenblatt for his perceptron. It was in 1971 that Paul Werbos, in his PhD thesis (P. J. Werbos 1974) introduced the BP algorithm, which made it possible to adjust the weights of a multilayer perceptron.

1.2.1 The Perceptron

From a conceptual point of view, a perceptron is a function made to perform a binary classification. In order to describe this function, let us first introduce the necessary notation and concepts associated with it. Assume a given observation that consists of a collection of D magnitudes observed from a physical phenomenon. These magnitudes are stored in a column vector, which will be called $\mathbf{x} \in \mathbb{R}^D$, which lies in a space of D dimensions. For illustrative purposes, let us construct a set of artificial data in a space of $D = 2$ dimensions as in Fig. 1.4.

The figure shows a set of points with coordinates $\mathbf{x} = (x_1, x_2)^\top$, where operator \top denotes the transpose operation, meaning that the vector is a column one even if it is written as a row vector. In this toy example, the data belongs to one of two classes (black or white) that we will label arbitrarily with the labels 1, -1 , though in some cases, labels 1, 0 are more convenient. It can be seen that the data is linearly separable, that is, both classes can be separated by placing a line between the black and white clusters of data. That is, roughly speaking, the idea of the perceptron. It must be trained to place a *separating hyperplane* between both classes. We define the hyperplane (particularized to a line in the two-dimensional example) as

$$\mathbf{w}^\top \mathbf{x} + b = 0 \tag{1.1}$$