
Make Your Own Neural Network

Python Machine Learning
Deep Learning and the Game of Go
The Nature of Code
Artificial Intelligence with Python
Introduction to Deep Learning
Neural Network Design
Machine Learning For Dummies
Deep Learning for Natural Language Processing
Neural Networks
Neural Network Projects with Python
Pattern Recognition and Neural Networks
Deep Learning for Coders with fastai and PyTorch
An Introduction to Neural Networks
Neural Networks and Deep Learning
Supervised Machine Learning for Text Analysis in
R
Multivariate Statistical Machine Learning Methods
for Genomic Prediction
Machine Learning with PyTorch and Scikit-Learn
Grokking Deep Learning
Deep Learning Illustrated
Modern Fortran
Deep Learning
Neural Networks for Beginners
Generative Deep Learning
Efficient Processing of Deep Neural Networks
Deep Learning with Python

Machine Learning with Neural Networks
Deep Learning from Scratch
Make Your Own Neural Network: An In-Depth
Visual Introduction for Beginners
Decision Trees and Random Forests
Neural Networks
Deep Learning For Dummies
Mathematics for Machine Learning
Make Your First GAN With PyTorch
Make Your Own Neural Network
Positive Intelligence
Artificial Neural Networks
The Math of Neural Networks
Neural Network for Beginners
Hands-On Neural Networks
Programming Machine Learning

*Make Your
Own Neural
Network*

*Downloaded
from
<ftp.wtvq.com>
by guest*

JUNE SIERRA

Python Machine Learning

Independently

Published

With the resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even

many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the details of important advanced architectures, implementing

everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically, computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural

networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework

Deep Learning and the Game of Go BPB Publications

A step-by-step visual journey through the mathematics of neural networks, and making your own using Python and Tensorflow. What you will gain from this book: * A deep understanding of how a Neural Network works. * How to build a Neural

Network from scratch using Python. Who this book is for: * Beginners who want to fully understand how networks work, and learn to build two step-by-step examples in Python. * Programmers who need an easy to read, but solid refresher, on the math of neural networks. What's Inside - 'Make Your Own Neural Network: An Indepth Visual Introduction For Beginners' What Is a Neural Network? Neural networks have made a gigantic comeback in the last few decades and you likely make use of them everyday without realizing it, but what exactly is a neural network? What is it used for and how does it fit within the broader arena of machine learning? we gently

explore these topics so that we can be prepared to dive deep further on. To start, we'll begin with a high-level overview of machine learning and then drill down into the specifics of a neural network. The Math of Neural Networks On a high level, a network learns just like we do, through trial and error. This is true regardless if the network is supervised, unsupervised, or semi-supervised. Once we dig a bit deeper though, we discover that a handful of mathematical functions play a major role in the trial and error process. It also becomes clear that a grasp of the underlying mathematics helps clarify how a network learns. * Forward Propagation *

Calculating The Total Error * Calculating The Gradients * Updating The Weights Make Your Own Artificial Neural Network: Hands on Example You will learn to build a simple neural network using all the concepts and functions we learned in the previous few chapters. Our example will be basic but hopefully very intuitive. Many examples available online are either hopelessly abstract or make use of the same data sets, which can be repetitive. Our goal is to be crystal clear and engaging, but with a touch of fun and uniqueness. This section contains the following eight chapters. Building Neural Networks in Python There are many ways to build a neural network and lots of

tools to get the job done. This is fantastic, but it can also be overwhelming when you start, because there are so many tools to choose from. We are going to take a look at what tools are needed and help you nail down the essentials. To build a neural network Tensorflow and Neural Networks There is no single way to build a feedforward neural network with Python, and that is especially true if you throw Tensorflow into the mix. However, there is a general framework that exists that can be divided into five steps and grouped into two parts. We are going to briefly explore these five steps so that we are prepared to use them to build a network later on.

Ready? Let's begin.

Neural Network:

Distinguish

Handwriting We are going to dig deep with Tensorflow and build a neural network that can distinguish between handwritten numbers. We'll use the same 5 steps we covered in the high-level overview, and we are going to take time exploring each line of code. Neural Network: Classify Images 10 minutes. That's all it takes to build an image classifier thanks to Google! We will provide a high-level overview of how to classify images using a convolutional neural network (CNN) and Google's Inception V3 model. Once finished, you will be able to tweak this code to classify any type of image sets! Cats, bats,

super heroes - the sky's the limit.

The Nature of Code

Greenleaf Book Group Python makes machine learning easy for beginners and experienced developers With computing power increasing exponentially and costs decreasing at the same time, there is no better time to learn machine learning using Python. Machine learning tasks that once required enormous processing power are now possible on desktop machines. However, machine learning is not for the faint of heart—it requires a good foundation in statistics, as well as programming knowledge. Python Machine Learning will help coders of all levels master one of the most

in-demand programming skillsets in use today. Readers will get started by following fundamental topics such as an introduction to Machine Learning and Data Science. For each learning algorithm, readers will use a real-life scenario to show how Python is used to solve the problem at hand. • Python data science—manipulating data and data visualization • Data cleansing • Understanding Machine learning algorithms • Supervised learning algorithms • Unsupervised learning algorithms • Deploying machine learning models Python Machine Learning is essential reading for students, developers, or anyone with a keen interest in taking their

coding skills to the next level.

[Artificial Intelligence with Python](#) "O'Reilly Media, Inc."

Text data is important for many domains, from healthcare to marketing to the digital humanities, but specialized approaches are necessary to create features for machine learning from language. Supervised Machine Learning for Text Analysis in R explains how to preprocess text data for modeling, train models, and evaluate model performance using tools from the tidyverse and tidymodels ecosystem. Models like these can be used to make predictions for new observations, to understand what natural language features or

characteristics contribute to differences in the output, and more. If you are already familiar with the basics of predictive modeling, use the comprehensive, detailed examples in this book to extend your skills to the domain of natural language processing. This book provides practical guidance and directly applicable knowledge for data scientists and analysts who want to integrate unstructured text data into their modeling pipelines. Learn how to use text data for both regression and classification tasks, and how to apply more straightforward algorithms like regularized regression or support vector machines as well as

deep learning approaches. Natural language must be dramatically transformed to be ready for computation, so we explore typical text preprocessing and feature engineering steps like tokenization and word embeddings from the ground up. These steps influence model results in ways we can measure, both in terms of model metrics and other tangible consequences such as how fair or appropriate model results are.

Introduction to Deep Learning MIT Press

An Introduction to Neural Networks falls into a new ecological niche for texts. Based on notes that have been class-tested for more than a decade, it is aimed at cognitive science and

neuroscience students who need to understand brain function in terms of computational modeling, and at engineers who want to go beyond formal algorithms to applications and computing strategies. It is the only current text to approach networks from a broad neuroscience and cognitive science perspective, with an emphasis on the biology and psychology behind the assumptions of the models, as well as on what the models might be used for. It describes the mathematical and computational tools needed and provides an account of the author's own ideas. Students learn how to teach arithmetic to a

neural network and get a short course on linear associative memory and adaptive maps. They are introduced to the author's brain-state-in-a-box (BSB) model and are provided with some of the neurobiological background necessary for a firm grasp of the general subject. The field now known as neural networks has split in recent years into two major groups, mirrored in the texts that are currently available: the engineers who are primarily interested in practical applications of the new adaptive, parallel computing technology, and the cognitive scientists and neuroscientists who are interested in scientific applications. As the gap between these two groups

widens, Anderson notes that the academics have tended to drift off into irrelevant, often excessively abstract research while the engineers have lost contact with the source of ideas in the field. Neuroscience, he points out, provides a rich and valuable source of ideas about data representation and setting up the data representation is the major part of neural network programming. Both cognitive science and neuroscience give insights into how this can be done effectively: cognitive science suggests what to compute and neuroscience suggests how to compute it. *Neural Network Design* Packt Publishing Ltd

There are many reasons why neural

networks fascinate us and have captivated headlines in recent years. They make web searches better, organize photos, and are even used in speech translation. Heck, they can even generate encryption. At the same time, they are also mysterious and mind-bending: how exactly do they accomplish these things? What goes on inside a neural network? On a high level, a network learns just like we do, through trial and error. This is true regardless if the network is supervised, unsupervised, or semi-supervised. Once we dig a bit deeper though, we discover that a handful of mathematical functions play a major role in the trial and error process. It also becomes clear

that a grasp of the underlying mathematics helps clarify how a network learns. In the following chapters we will unpack the mathematics that drive a neural network. To do this, we will use a feedforward network as our model and follow input as it moves through the network.

Machine Learning For Dummies

Independently

Published

Design and create neural networks with deep learning and artificial intelligence principles using OpenAI Gym, TensorFlow, and Keras
Key Features
Explore neural network architecture and understand how it functions
Learn algorithms to solve common problems

using back propagation and perceptrons
Understand how to apply neural networks to applications with the help of useful illustrations
Book Description
Neural networks play a very important role in deep learning and artificial intelligence (AI), with applications in a wide variety of domains, right from medical diagnosis, to financial forecasting, and even machine diagnostics.
Hands-On Neural Networks is designed to guide you through learning about neural networks in a practical way. The book will get you started by giving you a brief introduction to perceptron networks. You will then gain insights into machine learning and also understand what the future of AI could

look like. Next, you will study how embeddings can be used to process textual data and the role of long short-term memory networks (LSTMs) in helping you solve common natural language processing (NLP) problems. The later chapters will demonstrate how you can implement advanced concepts including transfer learning, generative adversarial networks (GANs), autoencoders, and reinforcement learning. Finally, you can look forward to further content on the latest advancements in the field of neural networks. By the end of this book, you will have the skills you need to build, train, and optimize your own neural network model that can be used to provide predictable

solutions. What you will learn

- Learn how to train a network by using backpropagation
- Discover how to load and transform images for use in neural networks
- Study how neural networks can be applied to a varied set of applications
- Solve common challenges faced in neural network development
- Understand the transfer learning concept to solve tasks using Keras and Visual Geometry Group (VGG) network
- Get up to speed with advanced and complex deep learning concepts like LSTMs and NLP
- Explore innovative algorithms like GANs and deep reinforcement learning

Who this book is for

If you are interested in artificial intelligence and deep learning and want to further your

skills, then this intermediate-level book is for you. Some knowledge of statistics will help you get the most out of this book.

Deep Learning for Natural Language Processing O'Reilly Media

If you want to learn how decision trees and random forests work, plus create your own, this visual book is for you. The fact is, decision tree and random forest algorithms are powerful and likely touch your life everyday. From online search to product development and credit scoring, both types of algorithms are at work behind the scenes in many modern applications and services. They are also used in countless industries such as

medicine, manufacturing and finance to help companies make better decisions and reduce risk. Whether coded or scratched out by hand, both algorithms are powerful tools that can make a significant impact. This book is a visual introduction for beginners that unpacks the fundamentals of decision trees and random forests. If you want to dig into the basics with a visual twist plus create your own algorithms in Python, this book is for you.

Neural Networks

Apress

The idea of simulating the brain was the goal of many pioneering works in Artificial Intelligence. The brain has been seen as a neural network, or a

set of nodes, or neurons, connected by communication lines. Currently, there has been increasing interest in the use of neural network models. This book contains chapters on basic concepts of artificial neural networks, recent connectionist architectures and several successful applications in various fields of knowledge, from assisted speech therapy to remote sensing of hydrological parameters, from fabric defect classification to application in civil engineering. This is a current book on Artificial Neural Networks and Applications, bringing recent advances in the area to the reader interested in this always-evolving

machine learning technique.

Neural Network Projects with Python

Springer

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the

obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify

data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

Pattern Recognition and Neural

Networks

Independently
Published

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve key metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or

increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as formalization and organization of key concepts from contemporary work

that provide insights that may spark new ideas.

[Deep Learning for Coders with fastai and PyTorch](#) Springer Science & Business Media

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show

you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering. Learn the latest deep learning techniques that matter most in practice. Improve accuracy, speed, and reliability by understanding how deep learning models work. Discover how to turn your models into web applications. Implement deep learning algorithms from scratch. Consider the ethical implications of your work. Gain insight from the

foreword by PyTorch cofounder, Soumith Chintala

An Introduction to Neural Networks

Addison-Wesley

Professional

All aboard The Coding Train! This beginner-friendly creative coding tutorial is designed to grow your skills in a fun, hands-on way as you build simulations of real-world phenomena with “The Coding Train” YouTube star Daniel Shiffman. What if you could recreate the awe-inspiring flocking patterns of birds or the hypnotic dance of fireflies—with code? For over a decade, The Nature of Code has empowered countless readers to do just that, bridging the gap between creative expression and programming. This

innovative guide by Daniel Shiffman, creator of the beloved Coding Train, welcomes budding and seasoned programmers alike into a world where code meets playful creativity. This JavaScript-based edition of Shiffman’s groundbreaking work gently unfolds the mysteries of the natural world, turning complex topics like genetic algorithms, physics-based simulations, and neural networks into accessible and visually stunning creations. Embark on this extraordinary adventure with projects involving: A physics engine: Simulate the push and pull of gravitational attraction. Flocking birds: Choreograph the

mesmerizing dance of a flock. Branching trees: Grow lifelike and organic tree structures. Neural networks: Craft intelligent systems that learn and adapt. Cellular automata: Uncover the magic of self-organizing patterns. Evolutionary algorithms: Play witness to natural selection in your code. Shiffman's work has transformed thousands of curious minds into creators, breaking down barriers between science, art, and technology, and inviting readers to see code not just as a tool for tasks but as a canvas for boundless creativity. Whether you're deciphering the elegant patterns of natural phenomena or crafting your own digital ecosystems, Shiffman's guidance is

sure to inform and inspire. The Nature of Code is not just about coding; it's about looking at the natural world in a new way and letting its wonders inspire your next creation. Dive in and discover the joy of turning code into art—all while mastering coding fundamentals along the way. NOTE: All examples are written with p5.js, a JavaScript library for creative coding, and are available on the book's website.

Neural Networks and Deep Learning John Wiley & Sons

Do you want to understand Neural Networks and learn everything about them but it looks like it is an exclusive club? Are you fascinated by Artificial Intelligence but you

think that it would be too difficult for you to learn? If you think that Neural Networks and Artificial Intelligence are the present and, even more, the future of technology, and you want to be part of it... well you are in the right place, and you are looking at the right book. If you are reading these lines you have probably already noticed this: Artificial Intelligence is all around you. Your smartphone that suggests you the next word you want to type, your Netflix account that recommends you the series you may like or Spotify's personalised playlists. This is how machines are learning from you in everyday life. And these examples are only the surface of this technological

revolution. Either if you want to start your own AI enterprise, to empower your business or to work in the greatest and most innovative companies, Artificial Intelligence is the future, and Neural Networks programming is the skill you want to have. The good news is that there is no exclusive club, you can easily (if you commit, of course) learn how to program and use neural networks, and to do that Neural Networks for Beginners is the perfect way. In this book you will learn: The types and components of neural networks The smartest way to approach neural network programming Why Algorithms are your friends The "three Vs" of Big Data (plus two new Vs) How machine learning will

help you making predictions The three most common problems with Neural Networks and how to overcome them Even if you don't know anything about programming, Neural Networks is the perfect place to start now. Still, if you already know about programming but not about how to do it in Artificial Intelligence, neural networks are the next thing you want to learn. And Neural Networks for Beginners is the best way to do it. Buy Neural Network for Beginners now to get the best start for your journey to Artificial Intelligence.

Supervised Machine Learning for Text Analysis in R Simon and Schuster
Build real-world Artificial Intelligence

applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are

looking to use Artificial Intelligence techniques in their existing technology stacks.

What You Will Learn

Realize different classification and regression techniques
 Understand the concept of clustering and how to use it to automatically segment data
 See how to build an intelligent recommender system
 Understand logic programming and how to use it
 Build automatic speech recognition systems
 Understand the basics of heuristic search and genetic programming
 Develop games using Artificial Intelligence
 Learn how reinforcement learning works
 Discover how to build intelligent applications centered on images, text, and time series data
 See

how to use deep learning algorithms and build applications based on it
 In Detail
 Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to

develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every

chapter, we explain an algorithm, implement it, and then build a smart application. [Multivariate Statistical Machine Learning Methods for Genomic Prediction](#) Pragmatic Bookshelf
Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to

learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book *Grokking Deep Learning* teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural

networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and

analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform.

Table of Contents

Introducing deep learning: why you should learn it

Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation

Introduction to neural learning: gradient descent

Learning multiple weights at a time: generalizing gradient descent

Building your first deep neural network: introduction to backpropagation

How to picture neural networks: in your head and on paper

Learning

signal and ignoring noise: introduction to regularization and batching

Modeling probabilities and nonlinearities: activation functions

Neural learning about edges and corners: intro to convolutional neural networks

Neural networks that understand language: king - man + woman == ?

Neural networks that write like Shakespeare: recurrent layers for variable-length data

Introducing automatic optimization: let's build a deep learning framework

Learning to write like Shakespeare: long short-term memory

Deep learning on unseen data: introducing federated learning

Where to go from here: a brief guide

Machine Learning with

PyTorch and Scikit-Learn O'Reilly Media

This modern and self-contained book offers a clear and accessible introduction to the important topic of machine learning with neural networks. In addition to describing the mathematical principles of the topic, and its historical evolution, strong connections are drawn with underlying methods from statistical physics and current applications within science and engineering. Closely based around a well-established undergraduate course, this pedagogical text provides a solid understanding of the key aspects of modern machine learning with artificial neural networks, for students in physics,

mathematics, and engineering. Numerous exercises expand and reinforce key concepts within the book and allow students to hone their programming skills. Frequent references to current research develop a detailed perspective on the state-of-the-art in machine learning research.

Grokking Deep

Learning CRC Press
Neural Networks presents concepts of neural-network models and techniques of parallel distributed processing in a three-step approach: - A brief overview of the neural structure of the brain and the history of neural-network modeling introduces to associative memory, preceptrons, feature-sensitive networks, learning strategies,

and practical applications. - The second part covers subjects like statistical physics of spin glasses, the mean-field theory of the Hopfield model, and the "space of interactions" approach to the storage capacity of neural networks. - The final part discusses nine programs with practical demonstrations of neural-network models. The software and source code in C are on a 3 1/2" MS-DOS diskette can be run with Microsoft, Borland, Turbo-C, or compatible compilers.

Deep Learning

Illustrated Simon and Schuster

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network

models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach

and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep

learning with NLP.

Modern Fortran

Springer

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided

throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the

concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.