

Build Your Own Neural Network Today With Step By Step Instructions Showing You How To Build Them Faster Than You Imagined Possible Using R

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch

Key Features

- Written by PyTorch's creator and key contributors
- Develop deep learning models in a familiar Pythonic way
- Use PyTorch to build an image classifier for cancer detection
- Diagnose problems with your neural network and improve training with data augmentation

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About The Book

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks.

What You Will Learn

- Understanding deep learning data structures such as tensors and neural networks
- Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results
- Implementing modules and loss functions
- Utilizing pretrained models from PyTorch Hub
- Methods for training networks with limited inputs
- Sifting through unreliable results to diagnose and fix problems in your neural network
- Improve your results with augmented data, better model architecture, and fine tuning

This Book Is Written For

For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required.

About The Authors

Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer.

Table of Contents

PART 1 - CORE PYTORCH

- 1 Introducing deep learning and the PyTorch Library
- 2 Pretrained networks
- 3 It starts with a tensor
- 4 Real-world data representation using tensors
- 5 The mechanics of learning
- 6 Using a neural network to fit the data
- 7 Telling birds from airplanes: Learning from images
- 8 Using convolutions to generalize

PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER

- 9 Using PyTorch to fight cancer
- 10 Combining data sources into a unified dataset
- 11 Training a classification model to detect suspected tumors
- 12 Improving training with metrics and augmentation
- 13 Using segmentation to find suspected nodules
- 14 End-to-end nodule analysis, and where to go next

PART 3 - DEPLOYMENT

- 15 Deploying to production

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 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 the 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 a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas.

Make Your Own Neural Network

Createspace Independent Publishing Platform

'Readers will emerge with a rigorous statistical grounding in the theory of how to construct and train neural networks in pattern recognition'

New Scientist

This book is a guide on how to implement a neural network in the Python programming language. It begins by giving you a brief overview of neural networks so as to know what they are, where they are used, and how they are implemented. The next step is an exploration of the backpropagation algorithm. This is the algorithm behind the functionality of neural networks, and it involves a forward and backward pass. NumPy is a Python library which can be used for the purpose of implementation of a neural network. This library is discussed in this book, and you are guided on how to use it for that purpose. The functionality of neural networks has to be improved. The various ways to improve how a neural network works is also explored. You are then guided on how to implement neural networks with Neupy, another Python library. The following topics are discussed in this book:

- A Brief Overview of Neural Networks
- Backpropagation Algorithm
- Neural Networks with Numpy
- Improving a Neural Network in Python
- Neupy - Models in Neural Networks

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

Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn:

- How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines
- How neural networks work and how they're trained
- How to use convolutional neural networks
- How to develop a successful deep learning model from scratch

You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, Practical Deep Learning will give you the skills and confidence to dive into your own machine learning projects.

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

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.

Discover best practices, reproducible architectures, and design patterns to help guide deep learning models from the lab into production. In *Deep Learning Patterns and Practices* you will learn: Internal functioning of modern convolutional neural networks Procedural reuse design pattern for CNN architectures Models for mobile and IoT devices Assembling large-scale model deployments Optimizing hyperparameter tuning Migrating a model to a production environment The big challenge of deep learning lies in taking cutting-edge technologies from R&D labs through to production. *Deep Learning Patterns and Practices* is here to help. This unique guide lays out the latest deep learning insights from author Andrew Ferlitsch's work with Google Cloud AI. In it, you'll find deep learning models presented in a unique new way: as extendable design patterns you can easily plug-and-play into your software projects. Each valuable technique is presented in a way that's easy to understand and filled with accessible diagrams and code samples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Discover best practices, design patterns, and reproducible architectures that will guide your deep learning projects from the lab into production. This awesome book collects and illuminates the most relevant insights from a decade of real world deep learning experience. You'll build your skills and confidence with each interesting example. About the book *Deep Learning Patterns and Practices* is a deep dive into building successful deep learning applications. You'll save hours of trial-and-error by applying proven patterns and practices to your own projects. Tested code samples, real-world examples, and a brilliant narrative style make even complex concepts simple and engaging. Along the way, you'll get tips for deploying, testing, and maintaining your projects. What's inside Modern convolutional neural networks Design pattern for CNN architectures Models for mobile and IoT devices Large-scale model deployments Examples for computer vision About the reader For machine learning engineers familiar with Python and deep learning. About the author Andrew Ferlitsch is an expert on computer vision, deep learning, and operationalizing ML in production at Google Cloud AI Developer Relations. Table of Contents PART 1 DEEP LEARNING FUNDAMENTALS 1 Designing modern machine learning 2 Deep neural networks 3 Convolutional and residual neural networks 4 Training fundamentals PART 2 BASIC DESIGN PATTERN 5 Procedural design pattern 6 Wide convolutional neural networks 7 Alternative connectivity patterns 8 Mobile convolutional neural networks 9 Autoencoders PART 3 WORKING WITH PIPELINES 10 Hyperparameter tuning 11 Transfer learning 12 Data distributions 13 Data pipeline 14 Training and deployment pipeline

Create and unleash the power of neural networks by implementing C# and .Net code Key Features Get a strong foundation of neural networks with access to various machine learning and deep learning libraries Real-world case studies illustrating various neural network techniques and architectures used by practitioners Cutting-edge coverage of Deep Networks, optimization algorithms, convolutional networks, autoencoders and many more Book Description Neural networks have made a surprise comeback in the last few years and have brought tremendous innovation in the world of artificial intelligence. The goal of this book is to provide C# programmers with practical guidance in solving complex computational challenges using neural networks and C# libraries such as CNTK, and TensorFlowSharp. This book will take you on a step-by-step practical journey, covering everything from the mathematical and theoretical aspects of neural networks, to building your own deep neural networks into your applications with the C# and .NET frameworks. This book begins by giving you a quick refresher of neural networks. You will learn how to build a neural network from scratch using packages such as Encog, Aforge, and Accord. You will learn about various concepts and techniques, such as deep networks, perceptrons, optimization algorithms, convolutional networks, and autoencoders. You will learn ways to add intelligent features to your .NET apps, such as facial and motion detection, object detection and labeling, language understanding, knowledge, and intelligent search. Throughout this book, you will be working on interesting demonstrations that will make it easier to implement complex neural networks in your enterprise applications. What you will learn Understand perceptrons and how to implement them in C# Learn how to train and visualize a neural network using cognitive services Perform image recognition for detecting and labeling objects using C# and TensorFlowSharp Detect specific image characteristics such as a face using Accord.Net Demonstrate particle swarm optimization using a simple XOR problem and Encog Train convolutional neural networks using ConvNetSharp Find optimal parameters for your neural network functions using numeric and heuristic optimization techniques. Who this book is for This book is for Machine Learning Engineers, Data Scientists, Deep Learning Aspirants and Data Analysts who are now looking to move into advanced machine learning and deep learning with C#. Prior knowledge of machine learning and working experience with C# programming is required to take most out of this book

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.

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

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.

Discover How to Build Your Own Neural Network From Scratch...Even if You've Got Zero Math or Coding Skills! What seemed like a lame and unbelievable sci-fi movie a few decades ago is now a reality. Machines can finally think. Maybe not quite as complex as the human brain, but more than enough to make everyone's life a lot easier. Artificial neural networks, based on the neurons found in the human brain give machines a 'brain'. Patterned just like biological neurons, these software or hardware are a variety of the deep learning technology. With their help you can make your computer learn by feeding it data, which will then be generated as the output you desire. It is they to thank for the nanoseconds in which computers operate. It may be science, but it is not actually rocket science. Everyone can learn how to take advantage of the progressed technology of today, get inside the 'brain' of the computers, and train them to perform the desired operations. They have been used in many different industries, and you can rest assured that you will find the perfect purpose for your own neural network. The best part about this book is that it doesn't require a college degree. Your high school math skills are quite enough for you to get a good grasp of the basics and learn how to build an artificial neural network. From non-mathematical explanations to teaching you the basic math behind the ANNs and training you how to actually program one, this book is the most helpful guide you will ever find. Carefully designed for you, the beginner, this guide will help you become a proud owner of a neural network in no time. Here's a Sneak Peak to What You'll Discover Inside this Book: The 6 unique benefits of neural networks The difference between biological and artificial neural networks And inside look into ANN (Artificial Neural Networks) The industries ANN is used in How to teach neural

networks to perform specific commands The different types of learning modalities (e.g. Hebbian Learning, unsupervised learning, supervised learning etc.) The architecture of ANN Basic math behind artificial neurons Simple networks for pattern classification The Hebb Rule How to build a simple neural network code The backpropagation algorithm and how to program it And much, much more! There's a lot more inside this book we'll cover, so be prepared. I've made to lucidly explain everything I cover so that there's zero confusion! Download this book today and discover all the intricate details of building your very own Neural Network

Deep learning neural networks have become easy to define and fit, but are still hard to configure. Discover exactly how to improve the performance of deep learning neural network models on your predictive modeling projects. With clear explanations, standard Python libraries, and step-by-step tutorial lessons, you'll discover how to better train your models, reduce overfitting, and make more accurate predictions.

Build your Machine Learning portfolio by creating 6 cutting-edge Artificial Intelligence projects using neural networks in Python Key Features Discover neural network architectures (like CNN and LSTM) that are driving recent advancements in AI Build expert neural networks in Python using popular libraries such as Keras Includes projects such as object detection, face identification, sentiment analysis, and more Book Description Neural networks are at the core of recent AI advances, providing some of the best resolutions to many real-world problems, including image recognition, medical diagnosis, text analysis, and more. This book goes through some basic neural network and deep learning concepts, as well as some popular libraries in Python for implementing them. It contains practical demonstrations of neural networks in domains such as fare prediction, image classification, sentiment analysis, and more. In each case, the book provides a problem statement, the specific neural network architecture required to tackle that problem, the reasoning behind the algorithm used, and the associated Python code to implement the solution from scratch. In the process, you will gain hands-on experience with using popular Python libraries such as Keras to build and train your own neural networks from scratch. By the end of this book, you will have mastered the different neural network architectures and created cutting-edge AI projects in Python that will immediately strengthen your machine learning portfolio. What you will learn Learn various neural network architectures and its advancements in AI Master deep learning in Python by building and training neural network Master neural networks for regression and classification Discover convolutional neural networks for image recognition Learn sentiment analysis on textual data using Long Short-Term Memory Build and train a highly accurate facial recognition security system Who this book is for This book is a perfect match for data scientists, machine learning engineers, and deep learning enthusiasts who wish to create practical neural network projects in Python. Readers should already have some basic knowledge of machine learning and neural networks.

One stop guide to implementing award-winning, and cutting-edge CNN architectures Key Features Fast-paced guide with use cases and real-world examples to get well versed with CNN techniques Implement CNN models on image classification, transfer learning, Object Detection, Instance Segmentation, GANs and more Implement powerful use-cases like image captioning, reinforcement learning for hard attention, and recurrent attention models Book Description Convolutional Neural Network (CNN) is revolutionizing several application domains such as visual recognition systems, self-driving cars, medical discoveries, innovative eCommerce and more. You will learn to create innovative solutions around image and video analytics to solve complex machine learning and computer vision related problems and implement real-life CNN models. This book starts with an overview of deep neural networks with the example of image classification and walks you through building your first CNN for human face detector. We will learn to use concepts like transfer learning with CNN, and Auto-Encoders to build very powerful models, even when not much of supervised training data of labeled images is available. Later we build upon the learning achieved to build advanced vision related algorithms for object detection, instance segmentation, generative adversarial networks, image captioning, attention mechanisms for vision, and recurrent models for vision. By the end of this book, you should be ready to implement advanced, effective and efficient CNN models at your professional project or personal initiatives by working on complex image and video datasets. What you will learn From CNN basic building blocks to advanced concepts understand practical areas they can be applied to Build an image classifier CNN model to understand how different components interact with each other, and then learn how to optimize it Learn different algorithms that can be applied to Object Detection, and Instance Segmentation Learn advanced concepts like attention mechanisms for CNN to improve prediction accuracy Understand transfer learning and implement award-winning CNN architectures like AlexNet, VGG, GoogLeNet, ResNet and more Understand the working of generative adversarial networks and how it can create new, unseen images Who this book is for This book is for data scientists, machine learning and deep learning practitioners, Cognitive and Artificial Intelligence enthusiasts who want to move one step further in building Convolutional Neural Networks. Get hands-on experience with extreme datasets and different CNN architectures to build efficient and smart ConvNet models. Basic knowledge of deep learning concepts and Python programming language is expected.

Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.

Tensorflow is the most popular Deep Learning Library out there. It has fantastic graph computations feature which helps data scientist to visualize his designed neural network using TensorBoard. This Machine learning library supports both Convolution as well as Recurrent Neural network. It supports parallel processing on CPU as well as GPU. Prominent machine learning algorithms supported by TensorFlow are Deep Learning Classification, wibe & deep, Boston Tree amongst others. The book is very hands-on and gives you industry ready deep learnings practices. Here is what is covered in the book

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of *Wait But Why* Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. *Deep Learning Illustrated* is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

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.

Convolutional Neural Networks in Python This book covers the basics behind Convolutional Neural Networks by introducing you to this complex world of deep learning and artificial neural networks in a simple and easy to understand way. It is perfect for any beginner out there looking forward to learning more about this machine learning field. This book is all about how to use convolutional neural networks for various image, object and other common classification problems in Python. Here, we also take a deeper look into various Keras layer used for building CNNs we take a look at different activation functions and much more, which will eventually lead you to creating highly accurate models able of performing great task results on various image classification, object classification and other problems. Therefore, at the end of the book, you will have a better insight into this world, thus you will be more than prepared to deal with more complex and challenging tasks on your own. Here Is a Preview of What You'll Learn In This Book... Convolutional neural networks structure How convolutional neural networks actually work Convolutional neural networks applications The importance of convolution operator Different convolutional neural networks layers and their importance Arrangement of spatial parameters How and when to use stride and zero-padding Method of parameter sharing Matrix multiplication and its importance Pooling and dense layers Introducing non-linearity relu activation function How to train your convolutional neural network models using backpropagation How and why to apply dropout CNN model training process How to build a convolutional neural network Generating predictions and calculating loss functions How to train and evaluate your MNIST classifier How to build a simple image classification CNN And much, much more!

Data assimilation is a process of fusing data with a model for the singular purpose of estimating unknown variables. It can be used, for example, to predict the evolution of the atmosphere at a given point and time. This book examines data assimilation methods including Kalman filtering, artificial intelligence, neural networks, machine learning, and cognitive computing.

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.

Starting with an easy introduction to KNIME Analytics Platform, this book will take you through the key features of the platform and cover the advanced and latest deep learning concepts in neural networks. In each chapter, you'll solve real-world case studies based on deep learning networks to spark your creativity for new projects.

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.

A gentle introduction to Generative Adversarial Networks, and a practical step-by-step tutorial on making your own with PyTorch. This beginner-friendly guide will give you hands-on experience: * understanding PyTorch basics * developing your first PyTorch neural network * exploring neural network refinements to improve performance * introduce CUDA GPU acceleration It will introduce GANs, one of the most exciting areas of machine learning: * introducing the concept step-by-step, in plain English * coding the simplest GAN to develop a good workflow * growing our confidence with an MNIST GAN * progressing to develop a GAN to generate full-colour human faces * experiencing how GANs fail, exploring remedies and improving GAN performance and stability Beyond the very basics, readers can explore more sophisticated GANs: * convolutional GANs for generated higher quality images * conditional GANs for generated images of a desired class The appendices will be useful for students of machine learning as they explain themes often skipped over in many courses: * calculating ideal loss values for balanced GANs * probability distributions and sampling them to create images * carefully chosen examples illustrating how convolutions work * a brief explanation of why gradient descent isn't suited to adversarial machine learning

??The Best Neural Networks Book for Beginners?? If you are looking for a complete beginners guide to learn neural networks with examples, in just a few hours, then you need to continue reading. Have you noticed the increasing prevalence of software that tries to learn from you? More and more, we are interacting with machines and platforms that try to predict what we are looking for. From movie and television show recommendations on Netflix based on your taste to the keyboard on your smartphone trying to predict and recommend the next word you may want to type, it's becoming obvious that machine learning will definitely be part of our future. If you are interested in learning more about the computer programs of tomorrow then, Understanding Neural Networks – A Practical Guide for Understanding and Programming Neural Networks and Useful Insights for Inspiring Reinvention is the book you have been waiting for. ?? Grab your copy today and learn ?? ? The history of neural networks and the way modern neural networks work ? How deep learning works ? The different types of neural networks ? The ability to explain a neural network to others, while simultaneously being able to build on this knowledge without being COMPLETELY LOST ? How to build your own neural network! ? An effective technique for hacking into a neural network ? Some introductory advice for modifying parameters in the code-based environment ? And much more... You'll be an Einstein in no time! And even if you are already up to speed on the topic, this book has the power to illustrate what a neural network is in a way that is capable of inspiring new approaches and technical improvements. The world can't wait to see what you can do! Most of all, this book will feed the abstract reasoning region of your mind so that you are able to theorize and invent new types and styles of machine learning. So, what are you waiting for? Scroll up and click the buy now button to learn everything you need to know in no time!

Tensorflow is the most popular Deep Learning Library out there. It has fantastic graph computations feature which helps data scientist to visualize his designed neural network using TensorBoard. This Machine learning library supports both Convolution as well as Recurrent Neural network. It supports parallel processing on CPU as well as GPU. Prominent machine learning algorithms supported by TensorFlow are Deep Learning Classification, wibe & deep, Boston Tree amongst others. The book is very hands-on and gives you industry ready deep learnings practices. Here is what is covered in the book – Table Of Content Chapter 1: What is Deep learning? Chapter 2: Machine Learning vs Deep Learning Chapter 3: What is TensorFlow? Chapter 4: Comparison of Deep Learning Libraries Chapter 5: How to Download and Install TensorFlow Windows and Mac Chapter 6: Jupyter Notebook Tutorial Chapter 7: Tensorflow on AWS Chapter 8: TensorFlow Basics: Tensor, Shape, Type, Graph, Sessions & Operators Chapter 9: Tensorboard: Graph Visualization with Example Chapter 10: NumPy Chapter 11: Pandas Chapter 12: Scikit-Learn Chapter 13: Linear Regression Chapter 14: Linear Regression Case Study Chapter 15: Linear Classifier in TensorFlow Chapter 16: Kernel Methods Chapter 17: TensorFlow ANN (Artificial Neural Network) Chapter 18: ConvNet(Convolutional Neural Network): TensorFlow Image Classification Chapter 19: Autoencoder with TensorFlow Chapter 20: RNN(Recurrent Neural Network) TensorFlow

Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/deep-?learning-with-r-in-motion). Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python as Deep Learning with Python by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image

Download Ebook Build Your Own Neural Network Today With Step By Step Instructions Showing You How To Build Them Faster Than You Imagined Possible Using R

classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions

Build your Own Neural Network today. Through easy-to-follow instruction and examples, you'll learn the fundamentals of Deep learning and build your very own Neural Network in Python using TensorFlow, Keras, PyTorch, and Theano. While you have the option of spending thousands of dollars on big and boring textbooks, we recommend getting the same pieces of information for a fraction of the cost. So Get Your Copy Now!! Why this book? Book Objectives The following are the objectives of this book: To help you understand deep learning in detail To help you know how to get started with deep learning in Python by setting up the coding environment. To help you transition from a deep learning Beginner to a Professional. To help you learn how to develop a complete and functional artificial neural network model in Python on your own. Who this Book is for? The author targets the following groups of people: Anybody who is a complete beginner to deep learning with Python. Anybody in need of advancing their Python for deep learning skills. Professors, lecturers or tutors who are looking to find better ways to explain Deep Learning to their students in the simplest and easiest way. Students and academicians, especially those focusing on python programming, neural networks, machine learning, and deep learning. What do you need for this Book? You are required to have installed the following on your computer: Python 3.X. TensorFlow . Keras . PyTorch The Author guides you on how to install the rest of the Python libraries that are required for deep learning. The author will guide you on how to install and configure the rest. What is inside the book? What is Deep Learning? An Overview of Artificial Neural Networks. Exploring the Libraries. Installation and Setup. TensorFlow Basics. Deep Learning with TensorFlow. Keras Basics. PyTorch Basics. Creating Convolutional Neural Networks with PyTorch. Creating Recurrent Neural Networks with PyTorch. From the back cover. Deep learning is part of machine learning methods based on learning data representations. This book written by Samuel Burns provides an excellent introduction to deep learning methods for computer vision applications. The author does not focus on too much math since this guide is designed for developers who are beginners in the field of deep learning. The book has been grouped into chapters, with each chapter exploring a different feature of the deep learning libraries that can be used in Python programming language. Each chapter features a unique Neural Network architecture including Convolutional Neural Networks. After reading this book, you will be able to build your own Neural Networks using Tensorflow, Keras, and PyTorch. Moreover, the author has provided Python codes, each code performing a different task. Corresponding explanations have also been provided alongside each piece of code to help the reader understand the meaning of the various lines of the code. In addition to this, screenshots showing the output that each code should return have been given. The author has used a simple language to make it easy even for beginners to understand.

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units BUILD YOUR OWN NEURAL NETWORK TODAY! With an EASY to follow process showing you how to build them FASTER than you imagined possible using R About This Book This rich, fascinating, accessible hands on guide, puts neural networks firmly into the hands of the practitioner. It reveals how they work, and takes you under the hood with an easy to follow process showing you how to build them faster than you imagined possible using the powerful, free R predictive analytics package. Everything you need to get started is contained within this book. It is your detailed, practical, tactical hands on guide. To accelerate your success, it contains exercises with fully worked solutions also provided. Once you have mastered the process, it will be easy for you to translate your knowledge into other powerful applications. A book for everyone interested in machine learning, predictive analytics, neural networks and decision science. Here is what it can do for you: SAVE TIME: Imagine having at your fingertips easy access to the very best neural network models without getting bogged down in mathematical details. In this book, you'll learn fast effective ways to build powerful neural network models easily using R. LEARN EASILY: Build Your Own Neural Network TODAY! Contains an easy to follow process showing you how to build the most successful neural networks used for learning from data; use this guide and build them easily and quickly. BOOST PRODUCTIVITY: Bestselling author and data scientist Dr. N.D. Lewis will show you how to build neural network models in less time than you ever imagined possible! Even if you're a busy professional, a student or hobbyist with little time, you will rapidly enhance your knowledge. EFFORTLESS SUCCESS: By spending as little as 10 minutes a day working through the dozens of real world examples, illustrations, practitioner tips and notes, you'll be able to make giant leaps forward in your knowledge, broaden your skill-set and generate new ideas for your own personal use. ELIMINATE ANXIETY: Forget trying to master every single mathematical detail, instead your goal is to simply to follow the process using real data that only takes about 5 to 15 minutes to complete. Within this process is a series of actions by which the neural network model is explained and constructed. All you have to do is follow the process. It is your checklist for use and reuse. 1 For people interested in statistics, machine learning, data analysis, data mining, and future hands-on practitioners seeking a career in the field, it sets a strong foundation, delivers the prerequisite knowledge, and whets your appetite for more. Here are some of the neural network models you will build: Multi layer Perceptrons Probabilistic Neural Networks Generalized Regression Neural Networks Recurrent Neural Networks Buy the book today. Your next big breakthrough using neural networks is only a page away!

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.

This text serves as a cookbook for neural network solutions to practical problems using C++. It will enable those with moderate programming experience to select a neural network model appropriate to solving a particular problem, and to produce a working program implementing that network. The book provides guidance along the entire problem-solving path, including designing the training set, preprocessing variables, training and validating the network, and evaluating its performance. Though the book is not intended as a general course in neural networks, no background in neural works is assumed and all models are presented from the ground up. The principle focus of the book is the three layer feedforward network, for more than a decade as the workhorse of professional arsenals. Other network models with strong performance records are also included. Bound in the book is an IBM diskette that includes the source code for all programs in the book. Much of this code can be easily adapted to C compilers. In addition, the operation of all programs is thoroughly discussed both in the text and in the comments within the code to facilitate translation to other languages.

A step-by-step gentle journey through the mathematics of neural networks, and making your own using the Python computer language. Neural networks are a key element of deep learning and artificial intelligence, which today is capable of some truly impressive feats. Yet too few really understand how neural networks actually work. This guide will take you on a fun and unhurried journey, starting from very simple ideas, and gradually building up an understanding of how neural networks work. You won't need any mathematics beyond secondary school, and an accessible introduction to calculus is also included. The ambition of this guide is to make neural networks as accessible as possible to as many readers as possible - there are enough texts for advanced readers already! You'll learn to code in Python and make your own neural network, teaching it to recognise human handwritten numbers, and performing as well as professionally developed networks. Part 1 is about ideas. We introduce the mathematical ideas underlying the neural networks, gently with lots of illustrations and examples. Part 2 is practical. We introduce the popular and easy to learn Python programming language, and gradually builds up a neural network which can learn to recognise human handwritten numbers, easily getting it to perform as well as networks made by professionals. Part 3 extends these ideas further. We push the performance of our neural network to an industry leading 98% using only simple ideas and code, test the network on your own handwriting, take a privileged peek inside the mysterious mind of a neural network, and even get it all working on a Raspberry Pi. All the code in this has been tested to work on a Raspberry Pi Zero.

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

[Copyright: 31c8c4f4521ba4dfd9a7ab855c9fdbe2](https://www.pdfdrive.com/build-your-own-neural-network-today-with-step-by-step-instructions-showing-you-how-to-build-them-faster-than-you-imagined-possible-using-r-ebook.html)