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

Hands-On Neural NetworksMake Your Own Neural NetworkThe PerceptronIntroduction to Deep LearningConvolutional Neural Networks in PythonNeural Network Programming With PythonPython Deep Learning: Develop Your First Neural Network in Python Using Tensorflow, Keras, and PytorchDeep Learning with PyTorchThe Nature of CodeJava Deep Learning CookbookMalware Data ScienceBetter Deep LearningMake Your Own Neural Network: An In-Depth Visual Introduction for BeginnersDeep Learning IllustratedHands-On Neural Network Programming with C#Strengthening Deep Neural NetworksLearning Deep Architectures for AIDeep LearningBrain On Fire: My Month of MadnessFrom Natural to Artificial IntelligenceNeural Networks and Deep LearningPractical FP in Scala (hard-Cover)Deep Learning for Coders with fastai and PyTorchNeural Networks for Pattern RecognitionNeural Network Projects with PythonHands-On Deep Learning Algorithms with PythonLearn TensorFlow in 24 HoursThe Math of Neural NetworksAn Introduction to Neural NetworksDeep Learning with PythonDeep Learning PipelineIntroduction to Information RetrievalDeep Learning with PyTorchPractical Computer Vision

Applications Using Deep Learning with CNNsBuild Your Own Neural Network Today!Deep Learning from ScratchDeep Learning With PythonMachine Learning With GoMake Your First GAN With PyTorchDeep Learning for Natural Language Processing

Hands-On Neural Networks

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

Make Your Own Neural Network

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, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary 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, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building

and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside Training deep neural networks Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI

tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. 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

The Perceptron

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

Page 6/47

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

Introduction to Deep Learning

This title shows you how to apply machine learning, statistics and data visualization as you build your own detection and intelligence system. Following an overview of basic reverse engineering concepts like static and dynamic analysis, you'll learn to measure code similarities in malware samples and use machine learning frameworks like scikit-learn and Keras to build and train your own detectors.

Convolutional Neural Networks in Python

Theoretical results suggest that in order to learn the kind of complicated functions that can represent high-

level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

Neural Network Programming With Python

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.

Python Deep Learning: Develop Your First Neural Network in Python Using Tensorflow, Keras, and Pytorch

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.

The Nature of Code

Class-tested and coherent, this textbook teaches classical and web information retrieval, including web search and the related areas of text classification and text clustering from basic concepts. It gives an up-to-date treatment of all aspects of the design and

implementation of systems for gathering, indexing, and searching documents; methods for evaluating systems; and an introduction to the use of machine learning methods on text collections. All the important ideas are explained using examples and figures, making it perfect for introductory courses in information retrieval for advanced undergraduates and graduate students in computer science. Based on feedback from extensive classroom experience, the book has been carefully structured in order to make teaching more natural and effective. Slides and additional exercises (with solutions for lecturers) are also available through the book's supporting website to help course instructors prepare their lectures.

Java Deep Learning Cookbook

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

Showing You How To Build Them Faster Than gained to build a question-answer chatbot system.
You Imagined Possible Using R 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.

Malware Data Science

"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.

Better Deep Learning

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual

Page 13/47

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.

Make Your Own Neural Network: An In-Depth Visual Introduction for Beginners

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.

Deep Learning Illustrated

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!

Hands-On Neural Network Programming with C#

Build your own pipeline based on modern TensorFlow approaches rather than outdated engineering concepts. This book shows you how to build a deep learning pipeline for real-life TensorFlow projects. You'll learn what a pipeline is and how it works so you can build a full application easily and rapidly. Then troubleshoot and overcome basic Tensorflow obstacles to easily create functional apps and deploy well-trained models. Step-by-step and example-

oriented instructions help you understand each step of the deep learning pipeline while you apply the most straightforward and effective tools to demonstrative problems and datasets. You'll also develop a deep learning project by preparing data, choosing the model that fits that data, and debugging your model to get the best fit to data all using Tensorflow techniques. Enhance your skills by accessing some of the most powerful recent trends in data science. If you've ever considered building your own image or text-tagging solution or entering a Kaggle contest, Deep Learning Pipeline is for you! What You'll Learn Develop a deep learning project using data Study and apply various models to your data Debug and troubleshoot the proper model suited for your data Who This Book Is For Developers, analysts, and data scientists looking to add to or enhance their existing skills by accessing some of the most powerful recent trends in data science. Prior experience in Python or other TensorFlow related languages and mathematics would be helpful.

Strengthening Deep Neural Networks

Build neural network models in text, vision and advanced analytics using PyTorch Key Features Learn PyTorch for implementing cutting-edge deep learning algorithms. Train your neural networks for higher speed and flexibility and learn how to implement them in various scenarios; Cover various advanced neural network architecture such as ResNet, Inception, DenseNet and more with practical examples; Book Description Deep learning powers the

most intelligent systems in the world, such as Google Voice, Siri, and Alexa. Advancements in powerful hardware, such as GPUs, software frameworks such as PyTorch, Keras, Tensorflow, and CNTK along with the availability of big data have made it easier to implement solutions to problems in the areas of text, vision, and advanced analytics. This book will get you up and running with one of the most cutting-edge deep learning libraries—PyTorch. PyTorch is grabbing the attention of deep learning researchers and data science professionals due to its accessibility, efficiency and being more native to Python way of development. You'll start off by installing PyTorch, then quickly move on to learn various fundamental blocks that power modern deep learning. You will also learn how to use CNN, RNN, LSTM and other networks to solve real-world problems. This book explains the concepts of various state-of-the-art deep learning architectures, such as ResNet, DenseNet, Inception, and Seq2Seq, without diving deep into the math behind them. You will also learn about GPU computing during the course of the book. You will see how to train a model with PyTorch and dive into complex neural networks such as generative networks for producing text and images. By the end of the book, you'll be able to implement deep learning applications in PyTorch with ease. What you will learn Use PyTorch for GPU-accelerated tensor computations Build custom datasets and data loaders for images and test the models using torchvision and torchtext Build an image classifier by implementing CNN architectures using PyTorch Build systems that do text classification and language modeling using RNN, LSTM, and GRU Learn advanced CNN architectures such as ResNet,

Inception, Densenet, and learn how to use them for transfer learning Learn how to mix multiple models for a powerful ensemble model Generate new images using GAN's and generate artistic images using style transfer Who this book is for This book is for machine learning engineers, data analysts, data scientists interested in deep learning and are looking to explore implementing advanced algorithms in PyTorch. Some knowledge of machine learning is helpful but not a mandatory need. Working knowledge of Python programming is expected.

Learning Deep Architectures for AI

Deep Learning

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!

Brain On Fire: My Month of Madness

Use Java and Deeplearning4j to build robust, scalable, and highly accurate AI models from scratch Key Features Install and configure Deeplearning4j to implement deep learning models from scratch Explore recipes for developing, training, and fine-tuning your neural network models in Java Model neural networks using datasets containing images, text, and time-series data Book Description Java is one of the most widely used programming languages in the world. With this book, you will see how to perform deep learning using Deeplearning4j (DL4J) – the most popular Java library for training neural networks efficiently. This book starts by showing you how to install and configure Java and DL4J on your system. You will then gain insights into deep learning basics and use your knowledge to create a deep neural network for binary classification from scratch. As you progress, you will discover how to build a convolutional neural network (CNN) in DL4J, and understand how to construct numeric vectors from text. This deep learning book will also guide you through performing anomaly detection on

unsupervised data and help you set up neural networks in distributed systems effectively. In addition to this, you will learn how to import models from Keras and change the configuration in a pre-trained DL4J model. Finally, you will explore benchmarking in DL4J and optimize neural networks for optimal results. By the end of this book, you will have a clear understanding of how you can use DL4J to build robust deep learning applications in Java. What you will learn Perform data normalization and wrangling using DL4J Build deep neural networks using DL4J Implement CNNs to solve image classification problems Train autoencoders to solve anomaly detection problems using DL4J Perform benchmarking and optimization to improve your model's performance Implement reinforcement learning for real-world use cases using RL4J Leverage the capabilities of DL4J in distributed systems Who this book is for If you are a data scientist, machine learning developer, or a deep learning enthusiast who wants to implement deep learning models in Java, this book is for you. Basic understanding of Java programming as well as some experience with machine learning and neural networks is required to get the most out of this book.

From Natural to Artificial Intelligence

Deploy deep learning applications into production across multiple platforms. You will work on computer vision applications that use the convolutional neural network (CNN) deep learning model and Python. This book starts by explaining the traditional machine-

learning pipeline, where you will analyze an image dataset. Along the way you will cover artificial neural networks (ANNs), building one from scratch in Python, before optimizing it using genetic algorithms. For automating the process, the book highlights the limitations of traditional hand-crafted features for computer vision and why the CNN deep-learning model is the state-of-art solution. CNNs are discussed from scratch to demonstrate how they are different and more efficient than the fully connected ANN (FCNN). You will implement a CNN in Python to give you a full understanding of the model. After consolidating the basics, you will use TensorFlow to build a practical image-recognition model that you will deploy to a web server using Flask, making it accessible over the Internet. Using Kivy and NumPy, you will create cross-platform data science applications with low overheads. This book will help you apply deep learning and computer vision concepts from scratch, step-by-step from conception to production. What You Will Learn Understand how ANNs and CNNs work Create computer vision applications and CNNs from scratch using Python Follow a deep learning project from conception to production using TensorFlow Use NumPy with Kivy to build cross-platform data science applications Who This Book Is ForData scientists, machine learning and deep learning engineers, software developers.

Neural Networks and Deep Learning

Practical FP in Scala: A hands-on approach, is a book for intermediate to advanced Scala developers. Aimed

at those who understand functional effects, referential transparency and the benefits of functional programming to some extent but who are missing some pieces to put all these concepts together to build a large application in a time-constrained manner.Throughout the chapters we will design, architect and develop a complete stateful application serving an API via HTTP, accessing a database and dealing with cached data, using the best practices and best functional libraries available in the Cats ecosystem.You will also learn about common design patterns such as managing state, error handling and anti-patterns, all accompanied by clear examples. Furthermore, at the end of the book, we will dive into some advanced concepts such as MTL, Classy Optics and Typeclass derivation.

Practical FP in Scala (hard-Cover)

Deep Learning for Coders with fastai and PyTorch

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

Neural Networks for Pattern Recognition

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.

Neural Network Projects with Python

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

Hands-On Deep Learning Algorithms with Python

How can we capture the unpredictable evolutionary and emergent properties of nature in software? How can understanding the mathematical principles behind our physical world help us to create digital worlds? This book focuses on a range of programming

strategies and techniques behind computer
simulations of natural systems, from elementary
concepts in mathematics and physics to more
advanced algorithms that enable sophisticated visual
results. Readers will progress from building a basic
physics engine to creating intelligent moving objects
and complex systems, setting the foundation for
further experiments in generative design. Subjects
covered include forces, trigonometry, fractals, cellular
automata, self-organization, and genetic algorithms.
The book's examples are written in Processing, an
open-source language and development environment
built on top of the Java programming language. On
the book's website (http: //www.natureofcode.com),
the examples run in the browser via Processing's
JavaScript mode.

Learn TensorFlow in 24 Hours

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.

The Math of Neural Networks

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 accelerationIt 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 stabilityBeyond the very basics, readers can explore more sophisticated GANs: * convolutional GANs for generated higher quality images * conditional GANs for generated images of a desired classThe 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

An Introduction to Neural Networks

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

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.

Deep Learning with Python

Build simple, maintainable, and easy to deploy machine learning applications. About This Book Build simple, but powerful, machine learning applications that leverage Go's standard library along with popular Go packages. Learn the statistics, algorithms, and techniques needed to successfully implement machine learning in Go Understand when and how to integrate certain types of machine learning model in Go applications. Who This Book Is For This book is for Go developers who are familiar with the Go syntax and can develop, build, and run basic Go programs. If you want to explore the field of machine learning and you love Go, then this book is for you! Machine Learning with Go will give readers the practical skills to perform the most common machine learning tasks with Go. Familiarity with some statistics and math topics is necessary. What You Will Learn Learn about data gathering, organization, parsing, and cleaning. Explore matrices, linear algebra, statistics, and probability. See how to evaluate and validate models. Look at regression, classification, clustering. Learn about neural networks and deep learning Utilize times

Page 32/47

series models and anomaly detection. Get to grip with techniques for deploying and distributing analyses and models. Optimize machine learning workflow techniques In Detail The mission of this book is to turn readers into productive, innovative data analysts who leverage Go to build robust and valuable applications. To this end, the book clearly introduces the technical aspects of building predictive models in Go, but it also helps the reader understand how machine learning workflows are being applied in real-world scenarios. Machine Learning with Go shows readers how to be productive in machine learning while also producing applications that maintain a high level of integrity. It also gives readers patterns to overcome challenges that are often encountered when trying to integrate machine learning in an engineering organization. The readers will begin by gaining a solid understanding of how to gather, organize, and parse real-work data from a variety of sources. Readers will then develop a solid statistical toolkit that will allow them to quickly understand gain intuition about the content of a dataset. Finally, the readers will gain hands-on experience implementing essential machine learning techniques (regression, classification, clustering, and so on) with the relevant Go packages. Finally, the reader will have a solid machine learning mindset and a powerful Go toolkit of techniques, packages, and example implementations. Style and approach This book connects the fundamental, theoretical concepts behind Machine Learning to practical implementations using the Go programming language.

Deep Learning Pipeline

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 Fran�ois 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 Fran�ois 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 Fran�ois 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

Introduction to Information Retrieval

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 with PyTorch

'My first serious blackout marked the line between sanity and insanity. Though I would have moments of lucidity over the coming days and weeks, I would never again be the same person ' Susannah Cahalan was a happy, clever, healthy twenty-four-year old. Then one day she woke up in hospital, with no memory of what had happened or how she had got there. Within weeks, she would be transformed into someone unrecognizable, descending into a state of acute psychosis, undergoing rages and convulsions, hallucinating that her father had murdered his wife; that she could control time with her mind. Everything she had taken for granted about her life, and who she was, was wiped out. Brain on Fire is Susannah's story of her terrifying descent into madness and the desperate hunt for a diagnosis, as, after dozens of tests and scans, baffled doctors concluded she should be confined in a psychiatric ward. It is also the story of how one brilliant man, Syria-born Dr Najar, finally proved - using a simple pen and paper - that Susannah's psychotic behaviour was caused by a rare autoimmune disease attacking her brain. His diagnosis of this little-known condition, thought to

have been the real cause of devil-possessions through history, saved her life, and possibly the lives of many others. Cahalan takes readers inside this newly-discovered disease through the progress of her own harrowing journey, piecing it together using memories, journals, hospital videos and records. Written with passionate honesty and intelligence, Brain on Fire is a searingly personal yet universal book, which asks what happens when your identity is suddenly destroyed, and how you get it back. 'With eagle-eye precision and brutal honesty, Susannah Cahalan turns her journalistic gaze on herself as she bravely looks back on one of the most harrowing and unimaginable experiences one could ever face: the loss of mind, body and self. Brain on Fire is a mesmerizing story' -Mira Bartók, New York Times bestselling author of The Memory Palace Susannah Cahalan is a reporter on the New York Post, and the recipient of the 2010 Silurian Award of Excellence in Journalism for Feature Writing. Her writing has also appeared in the New York Times, and is frequently picked up by the Daily Mail, Gawker, Gothamist, AOL and Yahoo among other news aggregrator sites.

Practical Computer Vision Applications Using Deep Learning with CNNs

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 ObjectivesThe 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 Tenserflow, 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.

Build Your Own Neural Network Today!

Understand basic to advanced deep learning algorithms, the mathematical principles behind them, and their practical applications. Key Features Get up-to-speed with building your own neural networks from scratch Gain insights into the mathematical principles behind deep learning algorithms Implement popular deep learning algorithms such as CNNs, RNNs, and more using TensorFlow Book Description Deep

learning is one of the most popular domains in the AI space, allowing you to develop multi-layered models of varying complexities. This book introduces you to popular deep learning algorithms—from basic to advanced—and shows you how to implement them from scratch using TensorFlow. Throughout the book, you will gain insights into each algorithm, the mathematical principles behind it, and how to implement it in the best possible manner. The book starts by explaining how you can build your own neural networks, followed by introducing you to TensorFlow, the powerful Python-based library for machine learning and deep learning. Moving on, you will get up to speed with gradient descent variants, such as NAG, AMSGrad, AdaDelta, Adam, and Nadam. The book will then provide you with insights into RNNs and LSTM and how to generate song lyrics with RNN. Next, you will master the math for convolutional and capsule networks, widely used for image recognition tasks. Then you learn how machines understand the semantics of words and documents using CBOW, skip-gram, and PV-DM. Afterward, you will explore various GANs, including InfoGAN and LSGAN, and autoencoders, such as contractive autoencoders and VAE. By the end of this book, you will be equipped with all the skills you need to implement deep learning in your own projects. What you will learn Implement basic-to-advanced deep learning algorithms Master the mathematics behind deep learning algorithms Become familiar with gradient descent and its variants, such as AMSGrad, AdaDelta, Adam, and Nadam Implement recurrent networks, such as RNN, LSTM, GRU, and seq2seq models Understand how machines interpret images using

CNN and capsule networks Implement different types of generative adversarial network, such as CGAN, CycleGAN, and StackGAN Explore various types of autoencoder, such as Sparse autoencoders, DAE, CAE, and VAE Who this book is for If you are a machine learning engineer, data scientist, AI developer, or simply want to focus on neural networks and deep learning, this book is for you. Those who are completely new to deep learning, but have some experience in machine learning and Python programming, will also find the book very helpful.

Deep Learning from Scratch

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

Deep Learning With Python

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.

Machine Learning With Go

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. Numby 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

Make Your First GAN With PyTorch

As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn't trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this adversarial input, and methods for increasing AI robustness to these attacks. If you're a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input Investigate methods used to generate adversarial input capable of fooling DNNs Explore real-world scenarios and model the adversarial threat Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data Examine some ways in which AI might become better at mimicking human perception in years to come

Deep Learning for Natural Language

Processing

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, wipe & 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

ROMANCE ACTION & ADVENTURE MYSTERY & THRILLER BIOGRAPHIES & HISTORY CHILDREN'S YOUNG ADULT FANTASY HISTORICAL FICTION HORROR LITERARY FICTION NON-FICTION SCIENCE FICTION