

Deep Learning Coursera

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.

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. The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology The world is a noisy and uncertain place. Probabilistic deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover probabilistic models that can indicate possible outcomes Learn to use normalizing flows for modeling and generating complex distributions Use Bayesian neural networks to access the uncertainty in the model About the reader For experienced machine learning developers. About the author Oliver Dürr is a professor at the University of Applied Sciences in Konstanz, Germany. Beate Sick holds a chair for applied statistics at ZHAW and works as a researcher and lecturer at the University of Zurich. Elvis Murina is a data scientist. Table of Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction to probabilistic deep learning 2 Neural network architectures 3 Principles of curve fitting PART 2 - MAXIMUM LIKELIHOOD APPROACHES FOR PROBABILISTIC DL MODELS 4 Building loss functions with the likelihood approach 5 Probabilistic deep learning models with TensorFlow Probability 6 Probabilistic deep learning models in the wild PART 3 - BAYESIAN APPROACHES FOR PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian neural networks Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork

descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

Understand deep learning, the nuances of its different models, and where these models can be applied. The abundance of data and demand for superior products/services have driven the development of advanced computer science techniques, among them image and speech recognition. Introduction to Deep Learning Using R provides a theoretical and practical understanding of the models that perform these tasks by building upon the fundamentals of data science through machine learning and deep learning. This step-by-step guide will help you understand the disciplines so that you can apply the methodology in a variety of contexts. All examples are taught in the R statistical language, allowing students and professionals to implement these techniques using open source tools. What You'll Learn Understand the intuition and mathematics that power deep learning models Utilize various algorithms using the R programming language and its packages Use best practices for experimental design and variable selection Practice the methodology to approach and effectively solve problems as a data scientist Evaluate the effectiveness of algorithmic solutions and enhance their predictive power Who This Book Is For Students, researchers, and data scientists who are familiar with programming using R. This book also is also of use for those who wish to learn how to appropriately deploy these algorithms in applications where they would be most useful.

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

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

Clustering is an important technique for discovering relatively dense sub-regions or sub-spaces of a multi-dimension data distribution. Clustering has been used in information retrieval for many different purposes, such as query expansion, document grouping, document indexing, and visualization of search results. In this book, we address issues of clustering algorithms, evaluation methodologies, applications, and architectures for information retrieval. The first two chapters discuss clustering algorithms. The chapter from Baeza-Yates et al. describes a clustering method for a general metric space which is a common model of data relevant to information retrieval. The chapter by Guha, Rastogi, and Shim presents a survey as well as detailed discussion of two clustering algorithms: CURE and ROCK for numeric data and categorical data respectively. Evaluation methodologies are addressed in the next two chapters. Ertoz et al. demonstrate the use of text retrieval benchmarks, such as TRECS, to evaluate clustering algorithms. He et al. provide objective measures of clustering quality in their chapter. Applications of clustering methods to information retrieval is addressed in the next four chapters. Chu et al. and Noel et al. explore feature selection using word stems, phrases, and link associations for document clustering and indexing. Wen et al. and Sung et al. discuss applications of clustering to user queries and data cleansing. Finally, we consider the problem of designing architectures for information retrieval. Crichton, Hughes, and Kelly elaborate on the development of a scientific data system architecture for information retrieval.

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is

still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

Get an example-driven introduction to Amazon Web Services (AWS), the cloud computing platform that enables you deploy highly scalable, secure, and reliable solutions—without the need for expensive hardware or dozens of engineers. You'll learn the profoundly powerful concepts that underlie cloud computing, and then explore the core AWS offerings and their key attributes. Real-world examples demonstrate how these concepts are put into practice. The book presents several hands-on projects that demonstrate the progressive evolution of an AWS-hosted system as requirements for scalability, security, reliability, and cost are accounted for.

Pattern recognition is a scientific discipline that is becoming increasingly important in the age of automation and information handling and retrieval. Patter Recognition, 2e covers the entire spectrum of pattern recognition applications, from image analysis to speech recognition and communications. This book presents cutting-edge material on neural networks, - a set of linked microprocessors that can form associations and uses pattern recognition to "learn" -and enhances student motivation by approaching pattern recognition from the designer's point of view. A direct result of more than 10 years of teaching experience, the text was developed by the authors through use in their own classrooms.

*Approaches pattern recognition from the designer's point of view *New edition highlights latest developments in this growing field, including independent components and support vector machines, not available elsewhere *Supplemented by computer examples selected from applications of interest

Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart programs to access interesting datasets from other web sites, collect data from users of your own applications, and analyze and understand the data once you've found it. Programming Collective Intelligence takes you into the world of machine learning and statistics, and explains how to draw conclusions about user experience, marketing, personal tastes, and human behavior in general -- all from information that you and others collect every day. Each algorithm is described clearly and concisely with code that can immediately be used on your web site, blog, Wiki, or specialized application. This book explains:

Collaborative filtering techniques that enable online retailers to recommend products or media
Methods of clustering to detect groups of similar items in a large dataset
Search engine features -- crawlers, indexers, query engines, and the PageRank algorithm
Optimization algorithms that search millions of possible solutions to a problem and choose the best one
Bayesian filtering, used in spam filters for classifying documents based on word types and other features
Using decision trees not only to make predictions, but to model the way decisions are made
Predicting numerical values rather than classifications to build price models
Support vector machines to match people in online dating sites
Non-negative matrix factorization to find the independent features in a dataset
Evolving intelligence for problem solving -- how a computer develops its skill by improving its own code the more it plays a game
Each chapter includes exercises for extending the algorithms to make them more powerful. Go beyond simple database-backed applications and put the wealth of Internet data to work for you. "Bravo! I cannot think of a better way for a developer to first learn these algorithms and methods, nor can I think of a better way for me (an old AI dog) to reinvigorate my knowledge of the details." -- Dan Russell, Google
"Toby's book does a great job of breaking down the complex subject matter of machine-learning algorithms into practical, easy-to-understand examples that can be directly applied to analysis of social interaction across the Web today. If I had this book two years ago, it would have saved precious time going down some fruitless paths." -- Tim Wolters, CTO, Collective Intellect

Master machine learning with SAS Viya! Machine learning can feel intimidating for new practitioners. Machine Learning with SAS Viya provides everything you need to know to get started with machine learning in SAS Viya, including decision trees, neural networks, and support vector machines. The analytics life cycle is covered from data preparation and discovery to deployment. Working with open-source code? Machine Learning with SAS Viya has you covered – step-by-step instructions are given on how to use SAS Model Manager tools with open source. SAS Model Studio features are highlighted to show how to carry out machine learning in SAS Viya. Demonstrations, practice tasks, and quizzes are included to help sharpen your skills. In this book, you will learn about: Supervised and unsupervised machine learning
Data preparation and dealing with missing and unstructured data
Model building and selection
Improving and optimizing models
Model deployment and monitoring performance

A surprisingly simple way for students to master any subject--based on one of the world's most popular online courses and the bestselling book A Mind for Numbers A Mind for Numbers and its wildly popular online companion course "Learning How to Learn" have empowered more than two million learners of all ages from around the world to master subjects that they once struggled with. Fans often wish they'd discovered these learning strategies earlier and ask how they can help their kids master these skills as well. Now in this new book for kids and teens, the authors reveal how to

make the most of time spent studying. We all have the tools to learn what might not seem to come naturally to us at first--the secret is to understand how the brain works so we can unlock its power. This book explains:

- Why sometimes letting your mind wander is an important part of the learning process
- How to avoid "rut think" in order to think outside the box
- Why having a poor memory can be a good thing
- The value of metaphors in developing understanding
- A simple, yet powerful, way to stop procrastinating

Filled with illustrations, application questions, and exercises, this book makes learning easy and fun.

With this practical book, AI and machine learning practitioners will learn how to successfully build and deploy data science projects on Amazon Web Services. The Amazon AI and machine learning stack unifies data science, data engineering, and application development to help level up your skills. This guide shows you how to build and run pipelines in the cloud, then integrate the results into applications in minutes instead of days. Throughout the book, authors Chris Fregly and Antje Barth demonstrate how to reduce cost and improve performance. Apply the Amazon AI and ML stack to real-world use cases for natural language processing, computer vision, fraud detection, conversational devices, and more. Use automated machine learning to implement a specific subset of use cases with SageMaker Autopilot. Dive deep into the complete model development lifecycle for a BERT-based NLP use case including data ingestion, analysis, model training, and deployment. Tie everything together into a repeatable machine learning operations pipeline. Explore real-time ML, anomaly detection, and streaming analytics on data streams with Amazon Kinesis and Managed Streaming for Apache Kafka. Learn security best practices for data science projects and workflows including identity and access management, authentication, authorization, and more.

"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

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

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in

size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

Discover ways to implement various deep learning algorithms by leveraging Python and other technologies Key Features Learn deep learning models through several activities Begin with simple machine learning problems, and finish by building a complex system of your own Teach your machines to see by mastering the technologies required for image recognition Book Description Deep learning is rapidly becoming the most preferred way of solving data problems. This is thanks, in part, to its huge variety of mathematical algorithms and their ability to find patterns that are otherwise invisible to us. Deep Learning from the Basics begins with a fast-paced introduction to deep learning with Python, its definition, characteristics, and applications. You'll learn how to use the Python interpreter and the script files in your applications, and utilize NumPy and Matplotlib in your deep learning models. As you progress through the book, you'll discover backpropagation—an efficient way to calculate the gradients of weight parameters—and study multilayer perceptrons and their limitations, before, finally, implementing a three-layer neural network and calculating multidimensional arrays. By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning. What you will learn Use Python with minimum external sources to implement deep learning programs Study the various deep learning and neural network theories Learn how to determine learning coefficients and the initial values of weights Implement trends such as Batch Normalization, Dropout, and Adam Explore applications like automatic driving, image generation, and reinforcement learning Who this book is for Deep Learning from the Basics is designed for data scientists, data analysts, and developers who want to use deep learning techniques to develop efficient solutions. This book is ideal for those who want a deeper understanding as well as an overview of the technologies. Some working knowledge of Python is a must. Knowledge of NumPy and pandas will be beneficial, but not essential.

Machine learning has finally come of age. With H2O software, you can perform machine learning and data analysis using a simple open source framework that's easy to use, has a wide range of OS and language support, and scales for big data. This hands-on guide teaches

you how to use H2O with only minimal math and theory behind the learning algorithms. If you're familiar with R or Python, know a bit of statistics, and have some experience manipulating data, author Darren Cook will take you through H2O basics and help you conduct machine-learning experiments on different sample data sets. You'll explore several modern machine-learning techniques such as deep learning, random forests, unsupervised learning, and ensemble learning. Learn how to import, manipulate, and export data with H2O Explore key machine-learning concepts, such as cross-validation and validation data sets Work with three diverse data sets, including a regression, a multinomial classification, and a binomial classification Use H2O to analyze each sample data set with four supervised machine-learning algorithms Understand how cluster analysis and other unsupervised machine-learning algorithms work

Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide About This Book Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide Real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is intended for a general audience of people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. What You Will Learn Learn about machine learning landscapes along with the historical development and progress of deep learning Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x Access public datasets and utilize them using TensorFlow to load, process, and transform

data Use TensorFlow on real-world datasets, including images, text, and more Learn how to evaluate the performance of your deep learning models Using deep learning for scalable object detection and mobile computing Train machines quickly to learn from data by exploring reinforcement learning techniques Explore active areas of deep learning research and applications In Detail Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples, you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing. 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.

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.

This book introduces machine learning methods in finance. It presents a unified treatment of machine learning and various statistical and computational disciplines in quantitative finance, such as financial econometrics and discrete time stochastic control, with an emphasis on how theory and hypothesis tests inform the choice of algorithm for financial data modeling and decision making. With the trend towards increasing computational resources and larger datasets, machine learning has grown into an important skillset for the finance industry. This book is written for advanced graduate students and academics in financial econometrics, mathematical finance and applied statistics, in addition to quants and data scientists in the field of quantitative finance. Machine Learning in Finance: From Theory to Practice is divided into three parts, each part covering theory and applications. The first presents supervised learning for cross-sectional data from both a Bayesian and frequentist perspective. The more advanced material places a firm emphasis on neural networks, including deep learning, as well as Gaussian processes, with examples in investment management and derivative modeling. The second part presents supervised learning for time series data, arguably the most common data type used in finance with examples in trading, stochastic volatility and fixed income modeling. Finally, the third part presents reinforcement learning and its applications in trading, investment and wealth management. Python code examples are provided to support the readers' understanding of the methodologies and applications. The book also includes more than 80 mathematical and programming exercises, with worked solutions available to instructors. As a bridge to research in this emergent field, the final chapter presents the frontiers of machine learning in finance from a researcher's perspective, highlighting how many well-known concepts in statistical physics are likely to emerge as important methodologies for machine learning in finance.

Create learning experiences that transform not only learning, but life itself. Learn about, improve, and expand your world of learning. This hands-on companion to the runaway best-seller, Deep Learning: Engage the World Change the World, provides an essential roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. Loaded with tips, tools, protocols, and real-world examples, the easy-to-use guide has everything educators need to construct and drive meaningful deep learning experiences that give purpose, unleash student potential, and prepare students to become problem-solving change agents in a global society.

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. Introduction to Pattern Recognition: A Matlab Approach is an accompanying manual to Theodoridis/Koutroumbas' Pattern Recognition. It includes Matlab code of the most common methods and algorithms in the book, together with a descriptive summary and solved examples, and including real-life data sets in imaging and audio recognition. This text is designed for electronic engineering, computer science, computer engineering, biomedical engineering and applied mathematics students taking graduate courses on pattern recognition and machine learning as well as R&D engineers and university researchers in image and signal processing/analysis, and computer vision. Matlab code and descriptive summary of the most common methods and algorithms in Theodoridis/Koutroumbas, Pattern Recognition, Fourth

Edition Solved examples in Matlab, including real-life data sets in imaging and audio recognition Available separately or at a special package price with the main text (ISBN for package: 978-0-12-374491-3)

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

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

This volume comprises six well-versed contributed chapters devoted to report the latest findings on the applications of machine learning for big data analytics. Big data is a term for data sets that are so large or complex that traditional data processing application software is inadequate to deal with them. The possible challenges in this direction include capture, storage, analysis, data curation, search, sharing, transfer, visualization, querying, updating and information privacy. Big data analytics is the process of examining large and varied data sets - i.e., big data - to uncover hidden patterns, unknown correlations, market trends, customer preferences and other useful information that can help organizations make more-informed business decisions. This volume is intended to be used as a reference by undergraduate and post graduate students of the disciplines of computer science, electronics and telecommunication, information science and electrical engineering. THE SERIES: FRONTIERS IN COMPUTATIONAL INTELLIGENCE The series Frontiers In Computational Intelligence is envisioned to provide comprehensive coverage and understanding of cutting edge research in computational intelligence. It intends to augment the scholarly discourse on all topics relating to the advances in artificial life and machine learning in the form of metaheuristics, approximate reasoning, and robotics. Latest research findings are coupled with applications to varied domains of engineering and computer sciences. This field is steadily growing especially with the advent of novel machine learning algorithms being applied to different domains of engineering and technology. The series brings together leading researchers that intend to continue to advance the field and create a broad knowledge about the most recent research.

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems,

NLP, and GANs

Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world.

What you will learn Cover advanced and state-of-the-art neural network architectures Understand the theory and math behind neural networks Train DNNs and apply them to modern deep learning problems Use CNNs for object detection and image segmentation Implement generative adversarial networks (GANs) and variational autoencoders to generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks

Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

Book Description How will AI evolve and what major innovations are on the horizon? What will its impact be on the job market, economy, and society? What is the path toward human-level machine intelligence? What should we be concerned about as artificial intelligence advances? Architects of Intelligence contains a series of in-depth, one-to-one interviews where New York Times bestselling author, Martin Ford, uncovers the truth behind these questions from some of the brightest minds in the Artificial Intelligence community. Martin has wide-ranging conversations with twenty-three of the world's foremost researchers and entrepreneurs working in AI and robotics: Demis Hassabis (DeepMind), Ray Kurzweil (Google), Geoffrey Hinton (Univ. of Toronto and Google), Rodney Brooks (Rethink Robotics), Yann LeCun (Facebook) , Fei-Fei Li (Stanford and Google), Yoshua Bengio (Univ. of Montreal), Andrew Ng (AI Fund), Daphne Koller

(Stanford), Stuart Russell (UC Berkeley), Nick Bostrom (Univ. of Oxford), Barbara Grosz (Harvard), David Ferrucci (Elemental Cognition), James Manyika (McKinsey), Judea Pearl (UCLA), Josh Tenenbaum (MIT), Rana el Kaliouby (Affectiva), Daniela Rus (MIT), Jeff Dean (Google), Cynthia Breazeal (MIT), Oren Etzioni (Allen Institute for AI), Gary Marcus (NYU), and Bryan Johnson (Kernel). Martin Ford is a prominent futurist, and author of Financial Times Business Book of the Year, Rise of the Robots. He speaks at conferences and companies around the world on what AI and automation might mean for the future.

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 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

This book begins with an introduction to AI, followed by machine learning, deep learning, NLP, and reinforcement

learning. Readers will learn about machine learning classifiers such as logistic regression, k-NN, decision trees, random forests, and SVMs. Next, the book covers deep learning architectures such as CNNs, RNNs, LSTMs, and auto encoders. Keras-based code samples are included to supplement the theoretical discussion. In addition, this book contains appendices for Keras, TensorFlow 2, and Pandas. Features: Covers an introduction to programming concepts related to AI, machine learning, and deep learning Includes material on Keras, TensorFlow2 and Pandas

[Copyright: 58396b6b2acf9cfdc4dc18bba0ebd194](#)