

Deep Learning Coursera

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

"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.

Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. Summary Most machine learning systems that are deployed in the world today learn from human feedback. However, most machine learning courses focus almost exclusively on the algorithms, not the human-computer interaction part of the systems. This can leave a big knowledge gap for data scientists working in real-world machine learning, where data scientists spend more time on data management than on building algorithms. Human-in-the-Loop Machine Learning is a practical guide to optimizing the entire machine learning process, including techniques for annotation, active learning, transfer learning, and using machine learning to optimize every step of the process. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Machine learning applications perform better with human feedback. Keeping the right people in the loop improves the accuracy of models, reduces errors in data, lowers costs, and helps you ship models faster. About the book Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. You'll find best practices on selecting sample data for human feedback, quality control for human annotations, and designing annotation interfaces. You'll learn to create training data for labeling, object detection, and semantic segmentation, sequence labeling, and more. The book starts with the basics and progresses to advanced techniques like transfer learning and self-supervision within annotation workflows. What's inside Identifying the right training and evaluation data Finding and managing people to annotate data Selecting annotation quality control strategies Designing annotation interfaces to improve accuracy and efficiency About the author Robert (Munro) Monarch is a data scientist and engineer who has built machine learning data for companies such as Apple, Amazon, Google, and IBM. He holds a PhD from Stanford. Robert holds a PhD from Stanford focused on Human-in-the-Loop machine learning for healthcare and disaster response, and is a disaster response professional in addition to being a machine learning professional. A worked example throughout this text is classifying disaster-related messages from real disasters that Robert has helped respond to in the past. Table of Contents PART 1 - FIRST STEPS 1 Introduction to human-in-the-loop machine learning 2 Getting started with human-in-the-loop machine learning PART 2 - ACTIVE LEARNING 3 Uncertainty sampling 4 Diversity sampling 5 Advanced active learning 6 Applying active learning to different machine learning tasks PART 3 - ANNOTATION 7 Working with the people annotating your data 8 Quality control for data annotation 9 Advanced data annotation and augmentation 10 Annotation quality for different machine learning tasks PART 4 - HUMAN-COMPUTER INTERACTION FOR MACHINE LEARNING 11 Interfaces for data annotation 12 Human-in-the-loop machine learning products

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.

A Step-by-Step Guide to Learning and Implementing Deep Learning Models Using R

Learning How to Learn

Modern Robotics

Deep Learning Illustrated

Clustering and Information Retrieval

A Guide to Building Deep Learning Systems

Getting Started with AWS

This book presents a compilation of selected papers from the 17th IEEE International Conference on Machine Learning and Applications (IEEE ICMLA 2018), focusing on use of deep learning technology in application like game playing, medical applications, video analytics, regression/classification, object

detection/recognition and robotic control in industrial environments. It highlights novel ways of using deep neural networks to solve real-world problems, and also offers insights into deep learning architectures and algorithms, making it an essential reference guide for academic researchers, professionals, software engineers in industry, and innovative product developers.

Unlock deeper insights into Machine Leaning with this vital guide to cutting-edge predictive analytics About This Book Leverage Python's most powerful open-source libraries for deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning

systems and algorithms Ask – and answer – tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning – whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data – its unique capabilities as a language let you build sophisticated algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success. Python Machine Learning gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

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

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code

one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way

to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of

complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and

sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need:

The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

Computational Intelligence

Applied Machine Learning

A Visual, Interactive Guide to Artificial Intelligence

Deep Learning with R

Architects of Intelligence

Generative Deep Learning

A Short Course

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.

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and

general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You ' ll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures,

TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you ' ll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning

models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

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 upyour 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

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

An Introduction

Methods and Applications

Probabilistic Deep Learning

Data Science on AWS

Active Learning and Annotation for Human-centered AI

Deep Learning for Coders with fastai and PyTorch

From Coding to Deep Learning

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** **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

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 Deep learning algorithms have brought a revolution to the computer vision community by introducing non-traditional and efficient solutions to several image-related problems that had long remained unsolved or partially addressed. This book presents a collection of eleven chapters where each individual chapter explains the deep learning principles of a specific topic, introduces reviews of up-to-date techniques, and presents research findings to the computer vision community. The book covers a broad scope of topics in deep learning concepts and applications such as accelerating the convolutional neural network inference on field-programmable gate arrays, fire detection in surveillance applications, face recognition, action and activity recognition, semantic segmentation for autonomous driving, aerial imagery registration, robot vision, tumor detection, and skin lesion segmentation as well as skin melanoma classification. The content of this book has been organized such that each chapter can be read independently from the others. The book is a valuable companion for researchers, for postgraduate and possibly senior undergraduate students who are taking an advanced course in related topics, and for those who are interested in deep learning with applications in computer vision, image processing, and pattern recognition.

with Applications in R

Your stepping stone to penetration testing

Data Science Job: How to become a Data Scientist

Learning from Data

Deep Learning in Natural Language Processing

Deep Learning

A Beginner's Guide to Finding Clients, Making Money, and Building Your Web Development Empire

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

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

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.

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.

Designing Next-Generation Machine Intelligence Algorithms

Grokking Deep Learning

AI and Machine Learning for Coders

Powerful, Scalable Techniques for Deep Learning and AI

Deep Learning Applications

Artificial Intelligence, Machine Learning, and Deep Learning

Mathematics for Machine LearningCambridge University Press

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

Financial Times Best Books of the Year 2018 TechRepublic Top Books Every Techie Should Read 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. Meet the minds behind the AI superpowers as they discuss the science, business and ethics of modern artificial intelligence. Read James Manyika's thoughts on AI analytics, Geoffrey Hinton's breakthroughs in AI programming and development, and Rana el Kaliouby's insights into AI marketing. This AI book collects the opinions of the luminaries of the AI business, such as Stuart Russell (coauthor of the leading AI textbook), Rodney Brooks (a leader in AI robotics), Demis Hassabis (chess prodigy and mind behind AlphaGo), and Yoshua Bengio (leader in deep learning) to complete your AI education and give you an AI advantage in 2019 and the future.

Machine learning methods are now an important tool for scientists, researchers, engineers and students in a wide range of areas. This book is written for people who want to adopt and use the main tools of machine learning, but aren't necessarily going to want to be machine learning researchers. Intended for students in final year undergraduate or first year graduate computer science programs in machine learning, this textbook is a machine learning toolkit. Applied Machine Learning covers many topics for people who want to use machine learning processes to get things done, with a strong emphasis on using existing tools and packages, rather than writing one's own code. A companion to the author's Probability and Statistics for Computer Science, this book picks up where the earlier book left off (but also supplies a summary of probability that the reader can use). Emphasizing the usefulness of standard machinery from applied statistics, this textbook gives an overview of the major applied areas in learning, including coverage of: • classification using standard machinery (naive bayes; nearest neighbor; SVM) • clustering and vector quantization (largely as in PSCS) • PCA (largely as in PSCS) • variants of PCA (NIPALS; latent semantic analysis; canonical correlation analysis) • linear regression (largely as in PSCS) • generalized linear models including logistic regression • model selection with Lasso, elasticnet • robustness and m-estimators • Markov chains and HMM's (largely as in PSCS) • EM in fairly gory detail; long experience teaching this suggests one detailed example is required, which students hate; but once they've been through that, the next one is easy • simple graphical models (in the variational inference section) • classification with neural networks, with a particular emphasis onimage classification • autoencoding with neural networks • structure learning

Mathematics for Machine Learning

Teaching Machines to Paint, Write, Compose, and Play

Deep Learning and the Game of Go

Programming Machine Learning

The truth about AI from the people building it

Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers

Learn Ethical Hacking from Scratch

Computational Intelligence: An Introduction, Second Edition offers an in-depth exploration into the adaptive mechanisms that enable intelligent behaviour in complex and changing environments. The main focus of this text is centred on the computational modelling of biological and natural intelligent systems, encompassing swarm intelligence, fuzzy systems, artificial neural networks, artificial immune systems and evolutionary computation. Engelbrecht provides readers with a wide knowledge of Computational Intelligence (CI) paradigms and algorithms; inviting readers to implement and problem solve real-world, complex problems within the CI development framework. This implementation framework will enable readers to tackle new problems without any difficulty through a single Java class as part of the CI library. Key features of this second edition include: A tutorial, hands-on based presentation of the material. State-of-the-art coverage of the most recent developments in computational intelligence with more elaborate discussions on intelligence and artificial intelligence (AI). New discussion of Darwinian evolution versus Lamarckian evolution, also including swarm robotics, hybrid systems and artificial immune systems. A section on how to perform empirical studies; topics including statistical analysis of stochastic algorithms, and an open source library of CI algorithms. Tables, illustrations, graphs, examples, assignments, Java code implementing the algorithms, and a complete CI implementation and experimental framework. Computational Intelligence: An Introduction, Second Edition is essential reading for third and fourth year undergraduate and postgraduate students studying CI. The first edition has been prescribed by a number of overseas universities and is thus a valuable teaching tool. In addition, it will also be a useful resource for researchers in Computational Intelligence and Artificial Intelligence, as well as engineers, statisticians, operational researchers, and bioinformaticians with an interest in applying AI or CI to solve problems in their domains. Check out http://www.ci.cs.up.ac.za for examples, assignments and Java code implementing the algorithms.

We're living in a digital world. Most of our global economy is digital and the sheer volume of data is stupendous. It's 2020 and we're living in the future. Data Scientist is one of the hottest job on the market right now. Demand for data science is huge and will only grow, and it seems like it will grow much faster than the actual number of data scientists. So if you want to make a career change and become a data scientist, now is the time. This book will guide you through the process. From my experience of working with multiple companies as a project manager, a data science consultant or a CTO, I was able to see the process of hiring data scientists and building data science teams. I know what's important to land your first job as a data scientist, what skills you should acquire, what you should show during a job interview.

Learn how to hack systems like black hat hackers and secure them like security experts Key Features Understand how computer systems work and their vulnerabilities Exploit weaknesses and hack into machines to test their security Learn how to secure systems from hackers Book Description This book starts with the basics of ethical hacking, how to practice hacking safely and legally, and how to install and interact with Kali Linux and the Linux terminal. You will explore network hacking, where you will see how to test the security of wired and wireless networks. You'll also learn how to crack the password for any Wi-Fi network (whether it uses WEP, WPA, or WPA2) and spy on the connected devices. Moving on, you will discover how to gain access to remote computer systems using client-side and server-side attacks. You will also get the hang of post-exploitation techniques, including remotely controlling and interacting with the systems that you compromised. Towards the end of the book, you will be able to pick up web application hacking techniques. You'll see how to discover, exploit, and prevent a number of website vulnerabilities, such as XSS and

SQL injections. The attacks covered are practical techniques that work against real systems and are purely for educational purposes. At the end of each section, you will learn how to detect, prevent, and secure systems from these attacks. What you will learn Understand ethical hacking and the different fields and types of hackers Set up a penetration testing lab to practice safe and legal hacking Explore Linux basics, commands, and how to interact with the terminal Access password-protected networks and spy on connected clients Use server and client-side attacks to hack and control remote computers Control a hacked system remotely and use it to hack other systems Discover, exploit, and prevent a number of web application vulnerabilities such as XSS and SQL injections Who this book is for Learning Ethical Hacking from Scratch is for anyone interested in learning how to hack and test the security of systems like professional hackers and security experts.

An Introduction to Statistical Learning provides an accessible overview of the field of statistical learning, an essential toolset for making sense of the vast and complex data sets that have emerged in fields ranging from biology to finance to marketing to astrophysics in the past twenty years. This book presents some of the most important modeling and prediction techniques, along with relevant applications. Topics include linear regression, classification, resampling methods, shrinkage approaches, tree-based methods, support vector machines, clustering, and more. Color graphics and real-world examples are used to illustrate the methods presented. Since the goal of this textbook is to facilitate the use of these statistical learning techniques by practitioners in science, industry, and other fields, each chapter contains a tutorial on implementing the analyses and methods presented in R, an extremely popular open source statistical software platform. Two of the authors co-wrote The Elements of Statistical Learning (Hastie, Tibshirani and Friedman, 2nd edition 2009), a popular reference book for statistics and machine learning researchers. An Introduction to Statistical Learning covers many of the same topics, but at a level accessible to a much broader audience. This book is targeted at statisticians and non-statisticians alike who wish to use cutting-edge statistical learning techniques to analyze their data. The text assumes only a previous course in linear regression and no knowledge of matrix algebra.

Deep Learning in Computer Vision

Introduction to Deep Learning Using R

Principles and Applications

Python Machine Learning

With Python, Keras and TensorFlow Probability

Will College Pay Off?

Reinforcement Learning, second edition

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.

Generative modeling is one of the hottest topics in AI. It ' s now possible to teach a machine to excel at human endeavors such as painting, writing, and composing music. With this practical book, machine-learning engineers and data scientists will discover how to re-create some of the most impressive examples of generative deep learning models, such as variational autoencoders, generative adversarial networks (GANs), encoder-decoder models and world models. Author David Foster demonstrates the inner workings of each technique, starting with the basics of deep learning before advancing to some of the most cutting-edge algorithms in the field. Through tips and tricks, you ' ll understand how to make your models learn more efficiently and become more creative. Discover how variational autoencoders can change facial expressions in photos Build practical GAN examples from scratch, including CycleGAN for style transfer and MuseGAN for music generation Create recurrent generative models for text generation and learn how to improve the models using attention Understand how generative models can help agents to accomplish tasks within a reinforcement learning setting Explore the architecture of the Transformer (BERT, GPT-2) and image generation models such as ProGAN and StyleGAN

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. Ertöz 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.

The decision of whether to go to college, or where, is hampered by poor information and inadequate understanding of the financial risk involved. Adding to the confusion, the same degree can cost dramatically different amounts for different people. A barrage of advertising offers new degrees designed to lead to specific jobs, but we see no information on whether graduates ever get those jobs. Mix in a frenzied applications process, and pressure from politicians for "relevant" programs, and there is an urgent need to separate myth from reality. Peter Cappelli, an acclaimed expert in employment trends, the workforce, and education, provides hard evidence that counters conventional wisdom and helps us make cost-effective choices. Among the issues Cappelli analyzes are: What is the real link between a college degree and a job that enables you to pay off the cost of college, especially in a market that is in constant change? Why it may be a mistake to pursue degrees that will land you the hottest jobs because what is hot today is unlikely to be so by the time you graduate. Why the most expensive colleges may actually be the cheapest because of their ability to graduate students on time. How parents and students can find out what different colleges actually deliver to students and whether it is something that employers really want. College is the biggest expense for many families, larger even than the cost of the family home, and one that can bankrupt students and their parents if it works out poorly. Peter Cappelli offers vital insight for parents and students to make decisions that both make sense financially and provide the foundation that will help students make their way in the world.

TinyML

A Guide to the Most Important Financial Decision You'll Ever Make

How to Succeed in School Without Spending All Your Time Studying; A Guide for Kids and Teens

Image Classification, Object Detection, and Face Recognition in Python

Human-in-the-Loop Machine Learning

Learning TensorFlow

Practical Machine Learning with H2O

Are you ready to jump-start your freelance career? Freelance Newbie has you covered! In this book, you'll learn practical, actionable steps you can start using today to get your first client by the end of the week. Featuring all the methods, techniques, tips, tricks, and insights you need to succeed, Freelance Newbie was written by a working freelancer whose mission is to help people like you find personal success and financial independence. The material you'll read here has never been featured at a lower price — you simply cannot get this kind of value for less. We go through everything step-by-step with real-world examples so you know exactly what you need to do to become a successful freelancer. In Freelance Newbie, you'll learn how to:

- Develop a business plan from scratch
- Establish a suitable work environment
- Configure your own freelance website to generate quality leads
- Determine what services to offer (and what to do if you don't know how to do something)
- Figure out an appropriate pricing scheme for your services
- Find “starter” clients that pave the way for 5-star social proof and full-paying, long-term clients
- Draft effective proposals and contracts
- Advertise for free (or very, very cheaply)
- Deliver above-average customer service
- Efficiently complete client projects — time runs out FAST
- And much, much more!

This book can also be used as the perfect companion manual to the video course available on Udemy by RealToughCandy.

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

Provides an overview of general deep learning methodology and its applications to a variety of signal and information processing tasks

A modern and unified treatment of the mechanics, planning, and control of robots, suitable for a first course in robotics.

Deep Learning for Computer Vision

Machine Learning with TensorFlow, Second Edition

Fundamentals of Deep Learning

Freelance Newbie

An Introduction to Statistical Learning

Artificial Intelligence 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

Step-by-step tutorials on deep learning neural networks for computer vision in python with Keras.

In recent years, deep learning has fundamentally changed the landscapes of a number of areas in artificial intelligence, including speech, vision, natural language, robotics, and game playing. In particular, the striking success of deep learning in a wide variety of natural language processing (NLP) applications has served as a benchmark for the advances in one of the most important tasks in artificial intelligence. This book reviews the state of the art of deep learning research and its successful applications to major NLP tasks, including speech recognition and understanding, dialogue systems, lexical analysis, parsing, knowledge graphs, machine translation, question answering, sentiment analysis, social computing, and natural language generation from images. Outlining and analyzing various research frontiers of NLP in the deep learning era, it features self-contained, comprehensive chapters written by leading researchers in the field. A glossary of technical terms and commonly used acronyms in the intersection of deep learning and NLP is also provided. The book appeals to advanced undergraduate and graduate students, post-doctoral researchers, lecturers and industrial researchers, as well as anyone interested in deep learning and natural language processing.

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.