

Deep Learning With Python Machine Learning Mastery

Thank you for reading **Deep Learning With Python Machine Learning Mastery** . As you may know, people have look numerous times for their favorite novels like this Deep Learning With Python Machine Learning Mastery , but end up in infectious downloads.

Rather than enjoying a good book with a cup of tea in the afternoon, instead they juggled with some infectious bugs inside their desktop computer.

Deep Learning With Python Machine Learning Mastery is available in our digital library an online access to it is set as public so you can get it instantly.

Our books collection saves in multiple countries, allowing you to get the most less latency time to download any of our books like this one.

Kindly say, the Deep Learning With Python Machine Learning Mastery is universally compatible with any devices to read

Imbalanced Classification with Python - Jason Brownlee 2020-01-14

Imbalanced classification are those classification tasks where the distribution of examples across the classes is not equal. Cut through the equations, Greek letters, and confusion, and discover the specialized techniques data preparation techniques, learning algorithms, and performance metrics that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover how to confidently develop robust models for your own imbalanced classification projects.

Generative Adversarial Networks with Python - Jason Brownlee 2019-07-11

Step-by-step tutorials on generative adversarial networks in python for image synthesis and image translation.

Mastering Machine Learning Algorithms - Giuseppe Bonaccorso 2018-05-25

Explore and master the most important algorithms for solving complex machine learning problems. Key Features Discover high-performing machine learning algorithms and understand how they work in depth. One-stop solution to mastering supervised, unsupervised, and semi-supervised machine learning algorithms and their implementation. Master concepts related to algorithm tuning, parameter optimization, and more Book Description Machine learning is a subset of AI that aims to make modern-day computer systems smarter and more intelligent. The real power of machine learning resides in its algorithms, which make even the most difficult things capable of being handled by machines. However, with the advancement in the technology and requirements of data, machines will have to be smarter than they are today to meet the overwhelming data needs; mastering these algorithms and using them optimally is the need of the hour. Mastering Machine Learning Algorithms is your complete guide to quickly getting to grips with popular machine learning algorithms. You will be introduced to the most widely used algorithms in supervised, unsupervised, and semi-supervised machine learning, and will learn how to use them in the best possible manner. Ranging from Bayesian models to the MCMC algorithm to Hidden Markov models, this book will teach you how to extract features from your dataset and perform dimensionality reduction by making use of Python-based libraries such as scikit-learn. You will also learn how to use Keras and TensorFlow to train effective neural networks. If you are looking for a single resource to study, implement, and solve end-to-end machine learning problems and use-cases, this is the book you need. What you will learn Explore how a ML model can be trained, optimized, and evaluated Understand how to create and learn static and dynamic probabilistic models Successfully cluster high-dimensional data and evaluate model accuracy Discover how artificial neural networks work and how to train, optimize, and validate them Work with Autoencoders

and Generative Adversarial Networks Apply label spreading and propagation to large datasets Explore the most important Reinforcement Learning techniques Who this book is for This book is an ideal and relevant source of content for data science professionals who want to delve into complex machine learning algorithms, calibrate models, and improve the predictions of the trained model. A basic knowledge of machine learning is preferred to get the best out of this guide.

Deep Learning for Time Series Forecasting - Jason Brownlee 2018-08-30

Deep learning methods offer a lot of promise for time series forecasting, such as the automatic learning of temporal dependence and the automatic handling of temporal structures like trends and seasonality. With clear explanations, standard Python libraries, and step-by-step tutorial lessons you'll discover how to develop deep learning models for your own time series forecasting projects.

Mastering Machine Learning Algorithms - Giuseppe Bonaccorso 2020-01-31

Updated and revised second edition of the bestselling guide to exploring and mastering the most important algorithms for solving complex machine learning problems Key Features Updated to include new algorithms and techniques Code updated to Python 3.8 & TensorFlow 2.x New coverage of regression analysis, time series analysis, deep learning models, and cutting-edge applications Book Description Mastering Machine Learning Algorithms, Second Edition helps you harness the real power of machine learning algorithms in order to implement smarter ways of meeting today's overwhelming data needs. This newly updated and revised guide will help you master algorithms used widely in semi-supervised learning, reinforcement learning, supervised learning, and unsupervised learning domains. You will use all the modern libraries from the Python ecosystem – including NumPy and Keras – to extract features from varied complexities of data. Ranging from Bayesian models to the Markov chain Monte Carlo algorithm to Hidden Markov models, this machine learning book teaches you how to extract features from your dataset, perform complex dimensionality reduction, and train supervised and semi-supervised models by making use of Python-based libraries such as scikit-learn. You will also discover practical applications for complex techniques such as maximum likelihood estimation, Hebbian learning, and ensemble learning, and how to use TensorFlow 2.x to train effective deep neural networks. By the end of this book, you will be ready to implement and solve end-to-end machine learning problems and use case scenarios. What you will learn Understand the characteristics of a machine learning algorithm Implement algorithms from supervised, semi-supervised, unsupervised, and RL domains Learn how regression works in time-series analysis and risk prediction Create, model, and train complex probabilistic models Cluster high-dimensional data and evaluate model accuracy

Discover how artificial neural networks work – train, optimize, and validate them Work with autoencoders, Hebbian networks, and GANs Who this book is for This book is for data science professionals who want to delve into complex ML algorithms to understand how various machine learning models can be built. Knowledge of Python programming is required.

Machine Learning with Python - Abhishek Vijayvargia 2018-03-01

Providing code examples in python, this book introduces the concepts of machine learning with mathematical explanations and programming fundamentals. --

Python Machine Learning - Sebastian Raschka 2015-09-23

Unlock deeper insights into Machine Learning 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 datascience 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.

Statistical Methods for Machine Learning - Jason Brownlee 2018-05-30

Statistics is a pillar of machine learning. You cannot develop a deep understanding and application of machine learning without it. Cut through the equations, Greek letters, and confusion, and discover the topics in statistics that you need to know. Using clear explanations, standard Python libraries, and

step-by-step tutorial lessons, you will discover the importance of statistical methods to machine learning, summary stats, hypothesis testing, nonparametric stats, resampling methods, and much more.

Introduction to Time Series Forecasting With Python - Jason Brownlee 2017-02-16

Time series forecasting is different from other machine learning problems. The key difference is the fixed sequence of observations and the constraints and additional structure this provides. In this Ebook, finally cut through the math and specialized methods for time series forecasting. Using clear explanations, standard Python libraries and step-by-step tutorials you will discover how to load and prepare data, evaluate model skill, and implement forecasting models for time series data.

Mathematics for Machine Learning - Marc Peter Deisenroth 2020-04-23

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.

Feature Engineering for Machine Learning - Alice Zheng 2018-03-23

Feature engineering is a crucial step in the machine-learning pipeline, yet this topic is rarely examined on its own. With this practical book, you'll learn techniques for extracting and transforming features—the numeric representations of raw data—into formats for machine-learning models. Each chapter guides you through a single data problem, such as how to represent text or image data. Together, these examples illustrate the main principles of feature engineering. Rather than simply teach these principles, authors Alice Zheng and Amanda Casari focus on practical application with exercises throughout the book. The closing chapter brings everything together by tackling a real-world, structured dataset with several feature-engineering techniques. Python packages including numpy, Pandas, Scikit-learn, and Matplotlib are used in code examples. You'll examine: Feature engineering for numeric data: filtering, binning, scaling, log transforms, and power transforms Natural text techniques: bag-of-words, n-grams, and phrase detection Frequency-based filtering and feature scaling for eliminating uninformative features Encoding techniques of categorical variables, including feature hashing and bin-counting Model-based feature engineering with principal component analysis The concept of model stacking, using k-means as a featurization technique Image feature extraction with manual and deep-learning techniques

Long Short-Term Memory Networks With Python - Jason Brownlee 2017-07-20

The Long Short-Term Memory network, or LSTM for short, is a type of recurrent neural network that achieves state-of-the-art results on challenging prediction problems. In this laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about LSTMs. Using clear explanations, standard

Python libraries and step-by-step tutorial lessons you will discover what LSTMs are, and how to develop a suite of LSTM models to get the most out of the method on your sequence prediction problems.

Mastering PyTorch - Ashish Ranjan Jha 2021-02-12

Master advanced techniques and algorithms for deep learning with PyTorch using real-world examples
Key Features
Understand how to use PyTorch 1.x to build advanced neural network models
Learn to perform a wide range of tasks by implementing deep learning algorithms and techniques
Gain expertise in domains such as computer vision, NLP, Deep RL, Explainable AI, and much more
Book Description
Deep learning is driving the AI revolution, and PyTorch is making it easier than ever before for anyone to build deep learning applications. This PyTorch book will help you uncover expert techniques to get the most out of your data and build complex neural network models. The book starts with a quick overview of PyTorch and explores using convolutional neural network (CNN) architectures for image classification. You'll then work with recurrent neural network (RNN) architectures and transformers for sentiment analysis. As you advance, you'll apply deep learning across different domains, such as music, text, and image generation using generative models and explore the world of generative adversarial networks (GANs). You'll not only build and train your own deep reinforcement learning models in PyTorch but also deploy PyTorch models to production using expert tips and techniques. Finally, you'll get to grips with training large models efficiently in a distributed manner, searching neural architectures effectively with AutoML, and rapidly prototyping models using PyTorch and fast.ai. By the end of this PyTorch book, you'll be able to perform complex deep learning tasks using PyTorch to build smart artificial intelligence models. What you will learn
Implement text and music generating models using PyTorch
Build a deep Q-network (DQN) model in PyTorch
Export universal PyTorch models using Open Neural Network Exchange (ONNX)
Become well-versed with rapid prototyping using PyTorch with fast.ai
Perform neural architecture search effectively using AutoML
Easily interpret machine learning (ML) models written in PyTorch using Captum
Design ResNets, LSTMs, Transformers, and more using PyTorch
Find out how to use PyTorch for distributed training using the torch.distributed API
Who this book is for
This book is for data scientists, machine learning researchers, and deep learning practitioners looking to implement advanced deep learning paradigms using PyTorch 1.x. Working knowledge of deep learning with Python programming is required.

XGBoost With Python - Jason Brownlee 2016-08-05

XGBoost is the dominant technique for predictive modeling on regular data. The gradient boosting algorithm is the top technique on a wide range of predictive modeling problems, and XGBoost is the fastest implementation. When asked, the best machine learning competitors in the world recommend using XGBoost. In this Ebook, learn exactly how to get started and bring XGBoost to your own machine learning projects.

Deep Learning for Natural Language Processing - Jason Brownlee 2017-11-21

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.

Deep Learning - Josh Patterson 2017-07-28

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source Deeplearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular
Understand how deep networks evolved from neural network fundamentals
Explore the major deep network architectures, including Convolutional and Recurrent
Learn how to map specific deep networks to the right problem
Walk through the fundamentals of tuning general neural networks and specific deep network architectures
Use vectorization techniques for different data types with DataVec, DL4J's workflow tool
Learn how to use DL4J natively on Spark and Hadoop

Deep Learning With Python - Jason Brownlee 2016-05-13

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.

Deep Learning - Ian Goodfellow 2016-11-10

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.

Mastering Machine Learning with scikit-learn - Gavin Hackling 2017-07-24

Use scikit-learn to apply machine learning to real-world problems About This Book Master popular machine learning models including k-nearest neighbors, random forests, logistic regression, k-means, naive Bayes, and artificial neural networks Learn how to build and evaluate performance of efficient models using scikit-learn Practical guide to master your basics and learn from real life applications of machine learning Who This Book Is For This book is intended for software engineers who want to understand how common machine learning algorithms work and develop an intuition for how to use them, and for data scientists who want to learn about the scikit-learn API. Familiarity with machine learning fundamentals and Python are helpful, but not required. What You Will Learn Review fundamental concepts such as bias and variance Extract features from categorical variables, text, and images Predict the values of continuous variables using linear regression and K Nearest Neighbors Classify documents and images using logistic regression and support vector machines Create ensembles of estimators using bagging and boosting techniques Discover hidden structures in data using K-Means clustering Evaluate the performance of machine learning systems in common tasks In Detail Machine learning is the buzzword bringing computer science and statistics together to build smart and efficient models. Using powerful algorithms and techniques offered by machine learning you can automate any analytical model. This book examines a variety of machine learning models including popular machine learning algorithms such as k-nearest neighbors, logistic regression, naive Bayes, k-means, decision trees, and artificial neural networks. It discusses data preprocessing, hyperparameter optimization, and ensemble methods. You will build systems that classify documents, recognize images, detect ads, and more. You will learn to use scikit-learn's API to extract features from categorical variables, text and images; evaluate model performance, and develop an intuition for how to improve your model's performance. By the end of this book, you will master all required concepts of scikit-learn to build efficient models at work to carry out advanced tasks with the practical approach. Style and approach This book is motivated by the belief that you do not understand something until you can describe it simply. Work through toy problems to develop your understanding of the learning algorithms and models, then apply your learnings to real-life problems.

Thoughtful Machine Learning - Matthew Kirk 2014-09-26

Learn how to apply test-driven development (TDD) to machine-learning algorithms—and catch mistakes that could sink your analysis. In this practical guide, author Matthew Kirk takes you through the principles of TDD and machine learning, and shows you how to apply TDD to several machine-learning algorithms, including Naive Bayesian classifiers and Neural Networks. Machine-learning algorithms often have tests baked in, but they can't account for human errors in coding. Rather than blindly rely on machine-learning results as many researchers have, you can mitigate the risk of errors with TDD and write clean, stable machine-learning code. If you're familiar with Ruby 2.1, you're ready to start. Apply TDD to write and run tests before you start coding Learn the best uses and tradeoffs of eight machine learning algorithms Use real-world examples to test each algorithm through engaging, hands-on exercises Understand the similarities between TDD and the scientific method for validating solutions Be aware of the risks of machine learning, such as underfitting and overfitting data Explore techniques for improving your machine-learning models or data extraction

Probability for Machine Learning - Jason Brownlee 2019-09-24

Probability is the bedrock of machine learning. You cannot develop a deep understanding and application of machine learning without it. Cut through the equations, Greek letters, and confusion, and discover the topics in probability that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover the importance of probability to machine learning, Bayesian probability, entropy, density estimation, maximum likelihood, and much more.

Better Deep Learning - Jason Brownlee 2018-12-13

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.

Deep Learning for Computer Vision - Jason Brownlee 2019-04-04

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

Master Machine Learning Algorithms - Jason Brownlee 2016-03-04

You must understand the algorithms to get good (and be recognized as being good) at machine learning. In this Ebook, finally cut through the math and learn exactly how machine learning algorithms work, then implement them from scratch, step-by-step.

Optimization for Machine Learning - Jason Brownlee 2021-09-22

Optimization happens everywhere. Machine learning is one example of such and gradient descent is probably the most famous algorithm for performing optimization. Optimization means to find the best value of some function or model. That can be the maximum or the minimum according to some metric. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will learn how to find the optimum point to numerical functions confidently using modern optimization algorithms.

Programming Machine Learning - Paolo Perrotta 2020-03-31

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.

Mastering Reinforcement Learning with Python - Enes Bilgin 2020-12-18

Get hands-on experience in creating state-of-the-art reinforcement learning agents using TensorFlow and RLlib to solve complex real-world business and industry problems with the help of expert tips and best practices

Key Features

- Understand how large-scale state-of-the-art RL algorithms and approaches work
- Apply RL to solve complex problems in marketing, robotics, supply chain, finance, cybersecurity, and more
- Explore tips and best practices from experts that will enable you to overcome real-world RL challenges

Book Description

Reinforcement learning (RL) is a field of artificial intelligence (AI) used for creating self-learning autonomous agents. Building on a strong theoretical foundation, this book takes a practical approach and uses examples inspired by real-world industry problems to teach you about state-of-the-art RL. Starting with bandit problems, Markov decision processes, and dynamic programming, the book provides an in-depth review of the classical RL techniques, such as Monte Carlo methods and temporal-difference learning. After that, you will learn about deep Q-learning, policy gradient algorithms, actor-critic methods, model-based methods, and multi-agent reinforcement learning. Then, you'll be introduced to some of the key approaches behind the most successful RL implementations, such as domain randomization and curiosity-driven learning. As you advance, you'll explore many novel algorithms with advanced implementations using modern Python libraries such as TensorFlow and Ray's RLlib package. You'll also find out how to implement RL in areas such as robotics, supply chain management, marketing, finance, smart cities, and cybersecurity while assessing the trade-offs between different approaches and avoiding common pitfalls. By the end of this book, you'll have mastered how to train and deploy your own RL agents for solving RL problems. What you will learn

Model and solve complex sequential decision-making problems using RL

Develop a solid understanding of how state-of-the-art RL methods work

Use Python and TensorFlow to code RL algorithms from scratch

Parallelize and scale up your RL implementations using Ray's RLlib package

Get in-depth knowledge of a wide variety of RL topics

Understand the trade-offs between different RL approaches

Discover and address the challenges of implementing RL in the real world

Who this book is for

This book is for expert machine learning practitioners and researchers looking to focus on hands-on reinforcement learning with Python by implementing advanced deep reinforcement learning concepts in real-world projects. Reinforcement learning experts who want to advance their knowledge to tackle large-scale and complex sequential decision-making problems will also find this book useful. Working knowledge of Python programming and deep learning along with prior experience in reinforcement learning is required.

Machine Learning Algorithms From Scratch with Python - Jason Brownlee 2016-11-16

You must understand algorithms to get good at machine learning. The problem is that they are only ever explained using Math. No longer. In this Ebook, finally cut through the math and learn exactly how machine learning algorithms work. Using clear explanations, simple pure Python code (no libraries!) and step-by-step tutorials you will discover how to load and prepare data, evaluate model skill, and implement a suite of linear, nonlinear and ensemble machine learning algorithms from scratch.

Deep Belief Nets in C++ and CUDA C: Volume 1 - Timothy Masters 2018-04-23

Discover the essential building blocks of the most common forms of deep belief

networks. At each step this book provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. The first of three in a series on C++ and CUDA C deep learning and belief nets, **Deep Belief Nets in C++ and CUDA C: Volume 1** shows you how the structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a thought process that is capable of learning abstract concepts built from simpler primitives. As such, you'll see that a typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. All the routines and algorithms presented in the book are available in the code download, which also contains some libraries of related routines. What You Will Learn

- Employ deep learning using C++ and CUDA C
- Work with supervised feedforward networks
- Implement restricted Boltzmann machines
- Use generative samplings

Discover why these are important

Who This Book Is For

Those who have at least a basic knowledge of neural networks and some prior programming experience, although some C++ and CUDA C is recommended.

Fundamentals of Deep Learning - Nikhil Buduma 2017-05-25

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

Mastering Machine Learning for Penetration Testing - Chiheb Chebbi 2018-06-27

Become a master at penetration testing using machine learning with Python

Key Features

- Identify ambiguities and breach intelligent security systems
- Perform unique cyber attacks to breach robust systems
- Learn to leverage machine learning algorithms

Book Description

Cyber security is crucial for both businesses and individuals. As systems are getting smarter, we now see machine learning interrupting computer security. With the adoption of machine learning in upcoming security products, it's important for pentesters and security researchers to understand how these systems work, and to breach them for testing purposes. This book begins with the basics of machine learning and the algorithms used to build robust systems. Once you've gained a fair understanding of how security products leverage machine learning, you'll dive into the core concepts of breaching such systems. Through practical use cases, you'll see how to find loopholes and surpass a self-learning security system. As you make your way through the chapters, you'll focus on topics such as network intrusion detection and AV and IDS evasion. We'll also cover the best practices when identifying ambiguities, and extensive techniques to breach an intelligent system. By the end of this book, you will be well-versed with identifying loopholes in a self-learning security system and will be able to efficiently breach a machine learning system. What you will learn

Take

an in-depth look at machine learning Get to know natural language processing (NLP) Understand malware feature engineering Build generative adversarial networks using Python libraries Work on threat hunting with machine learning and the ELK stack Explore the best practices for machine learning Who this book is for This book is for pen testers and security professionals who are interested in learning techniques to break an intelligent security system. Basic knowledge of Python is needed, but no prior knowledge of machine learning is necessary.

Mastering Machine Learning with Python in Six Steps - Manohar Swamynathan 2019-10-01

Explore fundamental to advanced Python 3 topics in six steps, all designed to make you a worthy practitioner. This updated version's approach is based on the "six degrees of separation" theory, which states that everyone and everything is a maximum of six steps away and presents each topic in two parts: theoretical concepts and practical implementation using suitable Python 3 packages. You'll start with the fundamentals of Python 3 programming language, machine learning history, evolution, and the system development frameworks. Key data mining/analysis concepts, such as exploratory analysis, feature dimension reduction, regressions, time series forecasting and their efficient implementation in Scikit-learn are covered as well. You'll also learn commonly used model diagnostic and tuning techniques. These include optimal probability cutoff point for class creation, variance, bias, bagging, boosting, ensemble voting, grid search, random search, Bayesian optimization, and the noise reduction technique for IoT data. Finally, you'll review advanced text mining techniques, recommender systems, neural networks, deep learning, reinforcement learning techniques and their implementation. All the code presented in the book will be available in the form of iPython notebooks to enable you to try out these examples and extend them to your advantage. What You'll Learn Understand machine learning development and frameworks Assess model diagnosis and tuning in machine learning Examine text mining, natural language processing (NLP), and recommender systems Review reinforcement learning and CNN Who This Book Is For Python developers, data engineers, and machine learning engineers looking to expand their knowledge or career into machine learning area.

Learning TensorFlow - Tom Hope 2017-08-09

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

Machine Learning Mastery With Weka - Jason Brownlee 2016-06-23

Machine learning is not just for professors. Weka is a top machine learning platform that provides an easy-to-use graphical interface and state-of-the-art algorithms. In this Ebook, learn exactly how to get started with applied machine learning using the Weka platform.

Learning Deep Architectures for AI - Yoshua Bengio 2009

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

Data Preparation for Machine Learning - Jason Brownlee 2020-06-30

Data preparation involves transforming raw data in to a form that can be modeled using machine learning algorithms. Cut through the equations, Greek letters, and confusion, and discover the specialized data preparation techniques that you need to know to get the most out of your data on your next project. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover how to confidently and effectively prepare your data for predictive modeling with machine learning.

Machine Learning Mastery With Python - Jason Brownlee 2016-04-08

The Python ecosystem with scikit-learn and pandas is required for operational machine learning. Python is the rising platform for professional machine learning because you can use the same code to explore different models in R&D then deploy it directly to production. In this Ebook, learn exactly how to get started and apply machine learning using the Python ecosystem.

Basics of Linear Algebra for Machine Learning - Jason Brownlee 2018-01-24

Linear algebra is a pillar of machine learning. You cannot develop a deep understanding and application of machine learning without it. In this laser-focused Ebook, you will finally cut through the equations, Greek letters, and confusion, and discover the topics in linear algebra that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover what linear algebra is, the importance of linear algebra to machine learning, vector, and matrix operations, matrix factorization, principal component analysis, and much more.

Deep Learning with PyTorch - Luca Pietro Giovanni Antiga 2020-07-01

"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 with Python](#) - Francois Chollet 2017-11-30

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