

Deep Learning Microsoft

Right here, we have countless book **deep learning microsoft** and collections to check out. We additionally come up with the money for variant types and as well as type of the books to browse. The gratifying book, fiction, history, novel, scientific research, as with ease as various supplementary sorts of books are readily easy to get to here.

As this deep learning microsoft, it ends up being one of the favored books deep learning microsoft collections that we have. This is why you remain in the best website to look the unbelievable ebook to have.

Model-Based Machine Learning - Taylor & Francis Group 2018-12-07

The Atlas of AI Kate Crawford 2021-04-06
The hidden costs of artificial intelligence, from natural resources and labor to privacy and freedom What happens when artificial intelligence saturates political life and depletes

the planet? How is AI shaping our understanding of ourselves and our societies? In this book Kate Crawford reveals how this planetary network is fueling a shift toward undemocratic governance and increased inequality. Drawing on more than a decade of research, award-winning science, and technology, Crawford reveals how AI is a technology of extraction: from the energy and

minerals needed to build and sustain its infrastructure, to the exploited workers behind "automated" services, to the data AI collects from us. Rather than taking a narrow focus on code and algorithms, Crawford offers us a political and a material perspective on what it takes to make artificial intelligence and where it goes wrong. While technical systems present a veneer of objectivity, they are always systems of power. This is an urgent account of what is at stake as technology companies use artificial intelligence to reshape the world.

Large- Scale C++ Volume I John Lakos

2019-12-02

Writing reliable and maintainable C++ software is hard. Designing such software at scale adds a new set of challenges. Creating large-scale systems requires a practical understanding of logical design - beyond the theoretical concepts addressed in most popular texts. To be successful on an enterprise scale, developers must also address physical design, a dimension

of software engineering that may be unfamiliar even to expert developers. Drawing on over 30 years of hands-on experience building massive, mission-critical enterprise systems, John Lakos shows how to create and grow Software Capital. This groundbreaking volume lays the foundation for projects of all sizes and demonstrates the processes, methods, techniques, and tools needed for successful real-world, large-scale development. Up to date and with a solid engineering focus, *Large-Scale C++, Volume I: Process and Architecture*, demonstrates fundamental design concepts with concrete examples. Professional developers of all experience levels will gain insights that transform their approach to design and development by understanding how to Raise productivity by leveraging differences between infrastructure and application development Achieve exponential productivity gains through feedback and hierarchical reuse Embrace the component's role as the fundamental unit of

both logical and physical design Analyze how fundamental properties of compiling and linking affect component design Discover effective partitioning of logical content in appropriately sized physical aggregates Internalize the important differences among sufficient, complete, minimal, and primitive software Deliver solutions that simultaneously optimize encapsulation, stability, and performance Exploit the nine established levelization techniques to avoid cyclic physical dependencies Use lateral designs judiciously to avoid the “heaviness” of conventional layered architectures Employ appropriate architectural insulation techniques for eliminating compile-time coupling Master the multidimensional process of designing large systems using component-based methods This is the first of John Lakos’s three authoritative volumes on developing large-scale systems using C++. This book, written for fellow software practitioners, uses familiar C++ constructs to solve real-world problems while identifying (and

motivating) modern C++ alternatives. Together with the forthcoming Volume II: Design and Implementation and Volume III: Verification and Testing, Large-Scale C++ offers comprehensive guidance for all aspects of large-scale C++ software development. If you are an architect or project leader, this book will empower you to solve critically important problems right now – and serve as your go-to reference for years to come. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Interpretable Machine Learning Christoph Molnar 2020

This book is about making machine learning models and their decisions interpretable. After exploring the concepts of interpretability, you will learn about simple, interpretable models such as decision trees, decision rules and linear regression. Later chapters focus on general model-agnostic methods for interpreting black

box models like feature importance and accumulated local effects and explaining individual predictions with Shapley values and LIME. All interpretation methods are explained in depth and discussed critically. How do they work under the hood? What are their strengths and weaknesses? How can their outputs be interpreted? This book will enable you to select and correctly apply the interpretation method that is most suitable for your machine learning project.

Mistering Azure Machine Learning Christoph Korner 2022-05-10

Supercharge and automate your deployments to Azure Machine Learning clusters and Azure Kubernetes Service using Azure Machine Learning services Key Features: Implement end-to-end machine learning pipelines on Azure Train deep learning models using Azure compute infrastructure Deploy machine learning models using MLOps Book Description: Azure Machine Learning is a cloud service for accelerating and

managing the machine learning (ML) project life cycle that ML professionals, data scientists, and engineers can use in their day-to-day workflows. This book covers the end-to-end ML process using Microsoft Azure Machine Learning, including data preparation, performing and logging ML training runs, designing training and deployment pipelines, and managing these pipelines via MLOps. The first section shows you how to set up an Azure Machine Learning workspace; ingest and version datasets; as well as preprocess, label, and enrich these datasets for training. In the next two sections, you'll discover how to enrich and train ML models for embedding, classification, and regression. You'll explore advanced NLP techniques, traditional ML models such as boosted trees, modern deep neural networks, recommendation systems, reinforcement learning, and complex distributed ML training techniques - all using Azure Machine Learning. The last section will teach you how to deploy the trained models as a batch

pipeline or real-time scoring service using Docker, Azure Machine Learning clusters, Azure Kubernetes Services, and alternative deployment targets. By the end of this book, you'll be able to combine all the steps you've learned by building an MLOps pipeline. What You Will Learn: Understand the end-to-end ML pipeline Get to grips with the Azure Machine Learning workspace Ingest, analyze, and preprocess datasets for ML using the Azure cloud Train traditional and modern ML techniques efficiently using Azure ML Deploy ML models for batch and real-time scoring Understand model interoperability with ONNX Deploy ML models to FPGAs and Azure IoT Edge Build an automated MLOps pipeline using Azure DevOps Who this book is for: This book is for machine learning engineers, data scientists, and machine learning developers who want to use the Microsoft Azure cloud to manage their datasets and machine learning experiments and build an enterprise-grade ML architecture using

MLOps. This book will also help anyone interested in machine learning to explore important steps of the ML process and use Azure Machine Learning to support them, along with building powerful ML cloud applications. A basic understanding of Python and knowledge of machine learning are recommended.

Deep Learning - Li Deng 2014

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

Mastering .NET Machine Learning - Jamie Dixon 2016-03-29

Master the art of machine learning with .NET and gain insight into real-world applications

About This Book Based on .NET framework

4.6.1, includes examples on ASP.NET Core 1.0

Set up your business application to start using machine learning techniques Familiarize the

user with some of the more common .NET

libraries for machine learning Implement several common machine learning techniques Evaluate,

optimize and adjust machine learning models
Who This Book Is For This book is targeted at .Net developers who want to build complex machine learning systems. Some basic understanding of data science is required. What You Will Learn Write your own machine learning applications and experiments using the latest .NET framework, including .NET Core 1.0 Set up your business application to start using machine learning. Accurately predict the future using regressions. Discover hidden patterns using decision trees. Acquire, prepare, and combine datasets to drive insights. Optimize business throughput using Bayes Classifier. Discover (more) hidden patterns using KNN and Naive Bayes. Discover (even more) hidden patterns using K-Means and PCA. Use Neural Networks to improve business decision making while using the latest ASP.NET technologies. Explore “Big Data”, distributed computing, and how to deploy machine learning models to IoT devices - making machines self-learning and adapting

Along the way, learn about Open Data, Bing maps, and MBrace In Detail .Net is one of the widely used platforms for developing applications. With the meteoric rise of Machine learning, developers are now keen on finding out how can they make their .Net applications smarter. Also, .NET developers are interested into moving into the world of devices and how to apply machine learning techniques to, well, machines. This book is packed with real-world examples to easily use machine learning techniques in your business applications. You will begin with introduction to F# and prepare yourselves for machine learning using .NET framework. You will be writing a simple linear regression model using an example which predicts sales of a product. Forming a base with the regression model, you will start using machine learning libraries available in .NET framework such as Math.NET, Numl.NET and Accord.NET with the help of a sample application. You will then move on to writing

multiple linear regressions and logistic regressions. You will learn what is open data and the awesomeness of type providers. Next, you are going to address some of the issues that we have been glossing over so far and take a deep dive into obtaining, cleaning, and organizing our data. You will compare the utility of building a KNN and Naive Bayes model to achieve best possible results. Implementation of Kmeans and PCA using Accord.NET and Numl.NET libraries is covered with the help of an example application. We will then look at many of issues confronting creating real-world machine learning models like overfitting and how to combat them using confusion matrixes, scaling, normalization, and feature selection. You will now enter into the world of Neural Networks and move your line of business application to a hybrid scientific application. After you have covered all the above machine learning models, you will see how to deal with very large datasets using MBrace and how to deploy machine

learning models to Internet of Thing (IoT) devices so that the machine can learn and adapt on the fly Style and approach This book will guide you in learning everything about how to tackle the flood of data being encountered these days in your .NET applications with the help of popular machine learning libraries offered by the .NET framework.

Reprogramming The American Dream - Kevin Scott 2020-04-07

**** #1 Wall Street Journal Bestseller **** In this essential book written by a rural native and Silicon Valley veteran, Microsoft's Chief technology officer tackles one of the most critical issues facing society today: the future of artificial intelligence and how it can be realistically used to promote growth, even in a shifting employment landscape. There are two prevailing stories about AI: for heartland low- and middle-skill workers, a dystopian tale of steadily increasing job destruction; for urban knowledge workers and the professional class, a

utopian tale of enhanced productivity and convenience. But there is a third way to look at this technology that will revolutionize the workplace and ultimately the world. Kevin Scott argues that AI has the potential to create abundance and opportunity for everyone and help solve some of our most vexing problems. As the chief technology officer at Microsoft, he is deeply involved in the development of AI applications, yet mindful of their potential impact on workers—knowledge he gained firsthand growing up in rural Virginia. Yes, the AI Revolution will radically disrupt economics and employment for everyone for generations to come. But what if leaders prioritized the programming of both future technology and public policy to work together to find solutions ahead of the coming AI epoch? Like public health, the space program, climate change and public education, we need international understanding and collaboration on the future of AI and work. For Scott, the crucial question

facing all of us is this: How do we work to ensure that the continued development of AI allows us to keep the American Dream alive? In this thoughtful, informed guide, he offers a clear roadmap to find the answer.

Hands-On Machine Learning with Azure -

Thomas K Abraham 2018-10-31

Implement machine learning, cognitive services, and artificial intelligence solutions by leveraging Azure cloud technologies Key Features Learn advanced concepts in Azure ML and the Cortana Intelligence Suite architecture Explore ML Server using SQL Server and HDInsight capabilities Implement various tools in Azure to build and deploy machine learning models Book Description Implementing Machine learning (ML) and Artificial Intelligence (AI) in the cloud had not been possible earlier due to the lack of processing power and storage. However, Azure has created ML and AI services that are easy to implement in the cloud. Hands-On Machine Learning with Azure teaches you how to perform

advanced ML projects in the cloud in a cost-effective way. The book begins by covering the benefits of ML and AI in the cloud. You will then explore Microsoft's Team Data Science Process to establish a repeatable process for successful AI development and implementation. You will also gain an understanding of AI technologies available in Azure and the Cognitive Services APIs to integrate them into bot applications. This book lets you explore prebuilt templates with Azure Machine Learning Studio and build a model using canned algorithms that can be deployed as web services. The book then takes you through a preconfigured series of virtual machines in Azure targeted at AI development scenarios. You will get to grips with the ML Server and its capabilities in SQL and HDInsight. In the concluding chapters, you'll integrate patterns with other non-AI services in Azure. By the end of this book, you will be fully equipped to implement smart cognitive actions in your models. What you will learn

Discover the benefits of leveraging the cloud for ML and AI Use Cognitive Services APIs to build intelligent bots Build a model using canned algorithms from Microsoft and deploy it as a web service Deploy virtual machines in AI development scenarios Apply R, Python, SQL Server, and Spark in Azure Build and deploy deep learning solutions with CNTK, MMLSpark, and TensorFlow Implement model retraining in IoT, Streaming, and Blockchain solutions Explore best practices for integrating ML and AI functions with ADLA and logic apps Who this book is for If you are a data scientist or developer familiar with Azure ML and cognitive services and want to create smart models and make sense of data in the cloud, this book is for you. You'll also find this book useful if you want to bring powerful machine learning services into your cloud applications. Some experience with data manipulation and processing, using languages like SQL, Python, and R, will aid in understanding the concepts covered in this book

Deep Learning with Microsoft Cognitive Toolkit Quick Start Guide - Willem Meints

2019-02-27

Learn how to train popular deep learning architectures such as autoencoders, convolutional and recurrent neural networks while discovering how you can use deep learning models in your software applications with Microsoft Cognitive Toolkit

Key Features

Understand the fundamentals of Microsoft Cognitive Toolkit and set up the development environment
Train different types of neural networks using Cognitive Toolkit and deploy it to production
Evaluate the performance of your models and improve your deep learning skills

Book Description

Cognitive Toolkit is a very popular and recently open sourced deep learning toolkit by Microsoft. Cognitive Toolkit is used to train fast and effective deep learning models. This book will be a quick introduction to using Cognitive Toolkit and will teach you how to train and validate different types of neural

networks, such as convolutional and recurrent neural networks. This book will help you understand the basics of deep learning. You will learn how to use Microsoft Cognitive Toolkit to build deep learning models and discover what makes this framework unique so that you know when to use it. This book will be a quick, no-nonsense introduction to the library and will teach you how to train different types of neural networks, such as convolutional neural networks, recurrent neural networks, autoencoders, and more, using Cognitive Toolkit. Then we will look at two scenarios in which deep learning can be used to enhance human capabilities. The book will also demonstrate how to evaluate your models' performance to ensure it trains and runs smoothly and gives you the most accurate results. Finally, you will get a short overview of how Cognitive Toolkit fits in to a DevOps environment

What you will learn

Set up your deep learning environment for the Cognitive Toolkit on Windows and Linux Pre-

process and feed your data into neural networks
Use neural networks to make efficient
predictions and recommendations Train and
deploy efficient neural networks such as CNN
and RNN Detect problems in your neural
network using TensorBoard Integrate Cognitive
Toolkit with Azure ML Services for effective
deep learning Who this book is for Data
Scientists, Machine learning developers, AI
developers who wish to train and deploy
effective deep learning models using Microsoft
CNTK will find this book to be useful. Readers
need to have experience in Python or similar
object-oriented language like C# or Java.

Introduction to Multi-Armed Bandits

Aleksandrs Slivkins 2019-10-31

Multi-armed bandits is a rich, multi-disciplinary
area that has been studied since 1933, with a
surge of activity in the past 10-15 years. This is
the first book to provide a textbook like
treatment of the subject.

Pattern Recognition and Machine Learning -

Christopher M. Bishop 2016-08-23

This is the first textbook on pattern recognition
to present the Bayesian viewpoint. The book
presents approximate inference algorithms that
permit fast approximate answers in situations
where exact answers are not feasible. It uses
graphical models to describe probability
distributions when no other books apply
graphical models to machine learning. No
previous knowledge of pattern recognition or
machine learning concepts is assumed.

Familiarity with multivariate calculus and basic
linear algebra is required, and some experience
in the use of probabilities would be helpful
though not essential as the book includes a self-
contained introduction to basic probability
theory.

**An Introduction to Neural Information
Retrieval** - Bhaskar Mitra 2018-12-23

Efficient Query Processing for Scalable Web
Search will be a valuable reference for
researchers and developers working on This

tutorial provides an accessible, yet comprehensive, overview of the state-of-the-art of Neural Information Retrieval.

Foundations of Data Science Avrim Blum
2020-01-23

This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities,

analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

Automated Machine Learning with Microsoft Azure - Dennis Michael Sawyers
2021-04-23

A practical, step-by-step guide to using Microsoft's AutoML technology on the Azure Machine Learning service for developers and data scientists working with the Python programming language. Key Features: Create, deploy, productionalize, and scale automated machine learning solutions on Microsoft Azure. Improve the accuracy of your ML models through automatic data featurization and model training. Increase productivity in your

organization by using artificial intelligence to solve common problems. Book Description Automated Machine Learning with Microsoft Azure will teach you how to build high-performing, accurate machine learning models in record time. It will equip you with the knowledge and skills to easily harness the power of artificial intelligence and increase the productivity and profitability of your business. Guided user interfaces (GUIs) enable both novices and seasoned data scientists to easily train and deploy machine learning solutions to production. Using a careful, step-by-step approach, this book will teach you how to use Azure AutoML with a GUI as well as the AzureML Python software development kit (SDK). First, you'll learn how to prepare data, train models, and register them to your Azure Machine Learning workspace. You'll then discover how to take those models and use them to create both automated batch solutions using machine learning pipelines and real-time scoring

solutions using Azure Kubernetes Service (AKS). Finally, you will be able to use AutoML on your own data to not only train regression, classification, and forecasting models but also use them to solve a wide variety of business problems. By the end of this Azure book, you'll be able to show your business partners exactly how your ML models are making predictions through automatically generated charts and graphs, earning their trust and respect. What you will learn Understand how to train classification, regression, and forecasting ML algorithms with Azure AutoML Prepare data for Azure AutoML to ensure smooth model training and deployment Adjust AutoML configuration settings to make your models as accurate as possible Determine when to use a batch-scoring solution versus a real-time scoring solution Productionalize your AutoML and discover how to quickly deliver value Create real-time scoring solutions with AutoML and Azure Kubernetes Service Train a large number of

AutoML models at once using the AzureML Python SDK. Who this book is for: Data scientists, aspiring data scientists, machine learning engineers, or anyone interested in applying artificial intelligence or machine learning in their business will find this machine learning book useful. You need to have beginner-level knowledge of artificial intelligence and a technical background in computer science, statistics, or information technology before getting started. Familiarity with Python will help you implement the more advanced features found in the chapters, but even data analysts and SQL experts will be able to train ML models after finishing this book.

Microsoft Azure Essentials Azure Machine Learning - Jeff Barnes 2015-04-25

Microsoft Azure Essentials from Microsoft Press is a series of free ebooks designed to help you advance your technical skills with Microsoft Azure. This third ebook in the series introduces Microsoft Azure Machine Learning, a service

that a developer can use to build predictive analytics models (using training datasets from a variety of data sources) and then easily deploy those models for consumption as cloud web services. The ebook presents an overview of modern data science theory and principles, the associated workflow, and then covers some of the more common machine learning algorithms in use today. It builds a variety of predictive analytics models using real world data, evaluates several different machine learning algorithms and modeling strategies, and then deploys the finished models as machine learning web services on Azure within a matter of minutes. The ebook also expands on a working Azure Machine Learning predictive model example to explore the types of client and server applications you can create to consume Azure Machine Learning web services. Watch Microsoft Press's blog and Twitter (@MicrosoftPress) to learn about other free ebooks in the Microsoft Azure Essentials series.

Automatic Speech Recognition - Dong Yu

2014-11-11

This book provides a comprehensive overview of the recent advancement in the field of automatic speech recognition with a focus on deep learning models including deep neural networks and many of their variants. This is the first automatic speech recognition book dedicated to the deep learning approach. In addition to the rigorous mathematical treatment of the subject, the book also presents insights and theoretical foundation of a series of highly successful deep learning models.

Dual Learning - Tao Qin 2020-11-13

Many AI (and machine learning) tasks present in dual forms, e.g., English-to-Chinese translation vs. Chinese-to-English translation, speech recognition vs. speech synthesis, question answering vs. question generation, and image classification vs. image generation. Dual learning is a new learning framework that leverages the primal-dual structure of AI tasks to

obtain effective feedback or regularization signals in order to enhance the learning/inference process. Since it was first introduced four years ago, the concept has attracted considerable attention in multiple fields, and been proven effective in numerous applications, such as machine translation, image-to-image translation, speech synthesis and recognition, (visual) question answering and generation, image captioning and generation, and code summarization and generation. Offering a systematic and comprehensive overview of dual learning, this book enables interested researchers (both established and newcomers) and practitioners to gain a better understanding of the state of the art in the field. It also provides suggestions for further reading and tools to help readers advance the area. The book is divided into five parts. The first part gives a brief introduction to machine learning and deep learning. The second part introduces the algorithms based on the dual reconstruction

principle using machine translation, image translation, speech processing and other NLP/CV tasks as the demo applications. It covers algorithms, such as dual semi-supervised learning, dual unsupervised learning and multi-agent dual learning. In the context of image translation, it introduces algorithms including CycleGAN, DualGAN, DiscoGAN, cGAN and more recent techniques/applications. The third part presents various work based on the probability principle, including dual supervised learning and dual inference based on the joint-probability principle and dual semi-supervised learning based on the marginal-probability principle. The fourth part reviews various theoretical studies on dual learning and discusses its connections to other learning paradigms. The fifth part provides a summary and suggests future research directions.

Introducing Machine Learning Dino Esposito

2020-02-05

Master machine learning concepts and develop

real-world solutions Machine learning offers immense opportunities, and Introducing Machine Learning delivers practical knowledge to make the most of them. Dino and Francesco Esposito start with a quick overview of the foundations of artificial intelligence and the basic steps of any machine learning project. Next, they introduce Microsoft's powerful ML.NET library, including capabilities for data processing, training, and evaluation. They present families of algorithms that can be trained to solve real-life problems, as well as deep learning techniques utilizing neural networks. The authors conclude by introducing valuable runtime services available through the Azure cloud platform and consider the long-term business vision for machine learning. · 14-time Microsoft MVP Dino Esposito and Francesco Esposito help you · Explore what's known about how humans learn and how intelligent software is built · Discover which problems machine learning can address · Understand the machine

learning pipeline: the steps leading to a deliverable model · Use AutoML to automatically select the best pipeline for any problem and dataset · Master ML.NET, implement its pipeline, and apply its tasks and algorithms · Explore the mathematical foundations of machine learning · Make predictions, improve decision-making, and apply probabilistic methods · Group data via classification and clustering · Learn the fundamentals of deep learning, including neural network design · Leverage AI cloud services to build better real-world solutions faster About This Book · For professionals who want to build machine learning applications: both developers who need data science skills and data scientists who need relevant programming skills · Includes examples of machine learning coding scenarios built using the ML.NET library

[Hands-On Machine Learning with ML.NET](#) -

Jarred Capellman 2020-03-27

Create, train, and evaluate various machine

learning models such as regression, classification, and clustering using ML.NET, Entity Framework, and ASP.NET Core Key Features Get well-versed with the ML.NET framework and its components and APIs using practical examples Learn how to build, train, and evaluate popular machine learning algorithms with ML.NET offerings Extend your existing machine learning models by integrating with TensorFlow and other libraries Book Description Machine learning (ML) is widely used in many industries such as science, healthcare, and research and its popularity is only growing. In March 2018, Microsoft introduced ML.NET to help .NET enthusiasts in working with ML. With this book, you'll explore how to build ML.NET applications with the various ML models available using C# code. The book starts by giving you an overview of ML and the types of ML algorithms used, along with covering what ML.NET is and why you need it to build ML apps. You'll then explore the ML.NET

framework, its components, and APIs. The book will serve as a practical guide to helping you build smart apps using the ML.NET library. You'll gradually become well versed in how to implement ML algorithms such as regression, classification, and clustering with real-world examples and datasets. Each chapter will cover the practical implementation, showing you how to implement ML within .NET applications. You'll also learn to integrate TensorFlow in ML.NET applications. Later you'll discover how to store the regression model housing price prediction result to the database and display the real-time predicted results from the database on your web application using ASP.NET Core Blazor and SignalR. By the end of this book, you'll have learned how to confidently perform basic to advanced-level machine learning tasks in ML.NET. What you will learn

Understand the framework, components, and APIs of ML.NET using C#

Develop regression models using ML.NET for employee attrition and file

classification

Evaluate classification models for sentiment prediction of restaurant reviews

Work with clustering models for file type classifications

Use anomaly detection to find anomalies in both network traffic and login history

Work with ASP.NET Core Blazor to create an ML.NET enabled web application

Integrate pre-trained TensorFlow and ONNX models in a WPF ML.NET application for image classification and object detection

Who this book is for

If you are a .NET developer who wants to implement machine learning models using ML.NET, then this book is for you. This book will also be beneficial for data scientists and machine learning developers who are looking for effective tools to implement various machine learning algorithms. A basic understanding of C# or .NET is mandatory to grasp the concepts covered in this book effectively.

Deep Learning for Computer Vision

Rajalingappaa Shanmugamani 2018-01-23

Learn how to model and train advanced neural

networks to implement a variety of Computer Vision tasks

Key Features

- Train different kinds of deep learning model from scratch to solve specific problems in Computer Vision
- Combine the power of Python, Keras, and TensorFlow to build deep learning models for object detection, image classification, similarity learning, image captioning, and more
- Includes tips on optimizing and improving the performance of your models under various constraints

Book Description

Deep learning has shown its power in several application areas of Artificial Intelligence, especially in Computer Vision. Computer Vision is the science of understanding and manipulating images, and finds enormous applications in the areas of robotics, automation, and so on. This book will also show you, with practical examples, how to develop Computer Vision applications by leveraging the power of deep learning. In this book, you will learn different techniques related to object classification, object detection, image

segmentation, captioning, image generation, face analysis, and more. You will also explore their applications using popular Python libraries such as TensorFlow and Keras. This book will help you master state-of-the-art, deep learning algorithms and their implementation. What you will learn

- Set up an environment for deep learning with Python, TensorFlow, and Keras
- Define and train a model for image and video classification
- Use features from a pre-trained Convolutional Neural Network model for image retrieval
- Understand and implement object detection using the real-world Pedestrian Detection scenario
- Learn about various problems in image captioning and how to overcome them by training images and text together
- Implement similarity matching and train a model for face recognition
- Understand the concept of generative models and use them for image generation
- Deploy your deep learning models and optimize them for high performance

Who this book is for This book is targeted at

data scientists and Computer Vision practitioners who wish to apply the concepts of Deep Learning to overcome any problem related to Computer Vision. A basic knowledge of programming in Python—and some understanding of machine learning concepts—is required to get the best out of this book.

Righting Software Juval Löwy 2019-11-29

Right Your Software and Transform Your Career
Righting Software presents the proven, structured, and highly engineered approach to software design that renowned architect Juval Löwy has practiced and taught around the world. Although companies of every kind have successfully implemented his original design ideas across hundreds of systems, these insights have never before appeared in print. Based on first principles in software engineering and a comprehensive set of matching tools and techniques, Löwy's methodology integrates system design and project design. First, he describes the primary area where many software

architects fail and shows how to decompose a system into smaller building blocks or services, based on volatility. Next, he shows how to flow an effective project design from the system design; how to accurately calculate the project duration, cost, and risk; and how to devise multiple execution options. The method and principles in Righting Software apply regardless of your project and company size, technology, platform, or industry. Löwy starts the reader on a journey that addresses the critical challenges of software development today by righting software systems and projects as well as careers—and possibly the software industry as a whole. Software professionals, architects, project leads, or managers at any stage of their career will benefit greatly from this book, which provides guidance and knowledge that would otherwise take decades and many projects to acquire. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for

details.

Machine Learning with Microsoft Technologies -
Leila Etaati 2019-06-12

Know how to do machine learning with Microsoft technologies. This book teaches you to do predictive, descriptive, and prescriptive analyses with Microsoft Power BI, Azure Data Lake, SQL Server, Stream Analytics, Azure Databricks, HD Insight, and more. The ability to analyze massive amounts of real-time data and predict future behavior of an organization is critical to its long-term success. Data science, and more specifically machine learning (ML), is today's game changer and should be a key building block in every company's strategy. Managing a machine learning process from business understanding, data acquisition and cleaning, modeling, and deployment in each tool is a valuable skill set. Machine Learning with Microsoft Technologies is a demo-driven book that explains how to do machine learning with Microsoft technologies. You will gain valuable

insight into designing the best architecture for development, sharing, and deploying a machine learning solution. This book simplifies the process of choosing the right architecture and tools for doing machine learning based on your specific infrastructure needs and requirements. Detailed content is provided on the main algorithms for supervised and unsupervised machine learning and examples show ML practices using both R and Python languages, the main languages inside Microsoft technologies. What You'll Learn Choose the right Microsoft product for your machine learning solution Create and manage Microsoft's tool environments for development, testing, and production of a machine learning project Implement and deploy supervised and unsupervised learning in Microsoft products Set up Microsoft Power BI, Azure Data Lake, SQL Server, Stream Analytics, Azure Databricks, and HD Insight to perform machine learning Set up a data science virtual machine and test-drive

installed tools, such as Azure ML Workbench, Azure ML Server Developer, Anaconda Python, Jupyter Notebook, Power BI Desktop, Cognitive Services, machine learning and data analytics tools, and more Architect a machine learning solution factoring in all aspects of self service, enterprise, deployment, and sharing Who This Book Is For Data scientists, data analysts, developers, architects, and managers who want to leverage machine learning in their products, organization, and services, and make educated, cost-saving decisions about their ML architecture and tool set.

Microsoft Azure Machine Learning - Sumit Mund 2015-06-16

This book provides you with the skills necessary to get started with Azure Machine Learning to build predictive models as quickly as possible, in a very intuitive way, whether you are completely new to predictive analysis or an existing practitioner. The book starts by exploring ML Studio, the browser-based development

environment, and explores the first step—data exploration and visualization. You will then build different predictive models using both supervised and unsupervised algorithms, including a simple recommender system. The focus then shifts to learning how to deploy a model to production and publishing it as an API. The book ends with a couple of case studies using all the concepts and skills you have learned throughout the book to solve real-world problems.

Predictive Analytics with Microsoft Azure Machine Learning 2nd Edition Valentine Fontama 2015-08-19

Predictive Analytics with Microsoft Azure Machine Learning, Second Edition is a practical tutorial introduction to the field of data science and machine learning, with a focus on building and deploying predictive models. The book provides a thorough overview of the Microsoft Azure Machine Learning service released for general availability on February 18th, 2015 with

practical guidance for building recommenders, propensity models, and churn and predictive maintenance models. The authors use task oriented descriptions and concrete end-to-end examples to ensure that the reader can immediately begin using this new service. The book describes all aspects of the service from data ingress to applying machine learning, evaluating the models, and deploying them as web services. Learn how you can quickly build and deploy sophisticated predictive models with the new Azure Machine Learning from Microsoft. What's New in the Second Edition? Five new chapters have been added with practical detailed coverage of: Python Integration - a new feature announced February 2015 Data preparation and feature selection Data visualization with Power BI Recommendation engines Selling your models on Azure Marketplace
Learning Microsoft Cognitive Services
Larsen 2017-10-23

Learn to build interactive and efficient applications by leveraging 24 effective cognitive services APIs powered by Microsoft About This Book Explore the capabilities of 24 of the APIs released as part of the Cognitive Services platform Build intelligent apps that combine the power of computer vision, speech recognition, and language processing Give your apps human-like cognitive intelligence with this hands-on guide Who This Book Is For .NET developers who want to add AI capabilities to their applications will find this book useful. No knowledge of machine learning or AI is necessary to work through this book. What You Will Learn Identify a person through visual inspection and audio Reduce user effort by utilizing AI-like capabilities Understand how to analyze images and text in different ways Find out how to analyze images using Vision APIs Add video analysis to applications using Vision APIs Utilize Search to find anything you want Analyze text to extract information and explore text

structure In Detail Microsoft has revamped its Project Oxford to launch the all new Cognitive Services platform-a set of 30 APIs to add speech, vision, language, and knowledge capabilities to apps. This book will introduce you to 24 of the APIs released as part of Cognitive Services platform and show you how to leverage their capabilities. More importantly, you'll see how the power of these APIs can be combined to build real-world apps that have cognitive capabilities. The book is split into three sections: computer vision, speech recognition and language processing, and knowledge and search. You will be taken through the vision APIs at first as this is very visual, and not too complex. The next part revolves around speech and language, which are somewhat connected. The last part is about adding real-world intelligence to apps by connecting them to Knowledge and Search APIs. By the end of this book, you will be in a position to understand what Microsoft Cognitive Service can offer and how to use the different APIs. Style

and approach This book takes you through essential API capabilities and shows how to utilize them to suit the needs of your application.

[Learn Azure in a Month of Lunches, Second Edition](#) - Iain Foulds 2020-10-06

Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. Summary You can be incredibly productive with Azure without mastering every feature, function, and service. Learn Azure in a Month of Lunches, Second Edition gets you up and running quickly, teaching you the most important concepts and tasks in 21 practical bite-sized lessons. As you explore the examples, exercises, and labs, you'll pick up valuable skills immediately and take your first steps to Azure mastery! This fully revised new edition covers core changes to the Azure UI, new Azure

features, Azure containers, and the upgraded Azure Kubernetes Service. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microsoft Azure is vast and powerful, offering virtual servers, application templates, and prebuilt services for everything from data storage to AI. To navigate it all, you need a trustworthy guide. In this book, Microsoft engineer and Azure trainer Iain Foulds focuses on core skills for creating cloud-based applications. About the book Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. What's inside Understanding Azure beyond point-and-click Securing applications and data Automating your environment Azure services for machine learning, containers, and more About the reader

This book is for readers who can write and deploy simple web or client/server applications. About the author Iain Foulds is an engineer and senior content developer with Microsoft. Table of Contents PART 1 - AZURE CORE SERVICES 1 Before you begin 2 Creating a virtual machine 3 Azure Web Apps 4 Introduction to Azure Storage 5 Azure Networking basics PART 2 - HIGH AVAILABILITY AND SCALE 6 Azure Resource Manager 7 High availability and redundancy 8 Load-balancing applications 9 Applications that scale 10 Global databases with Cosmos DB 11 Managing network traffic and routing 12 Monitoring and troubleshooting PART 3 - SECURE BY DEFAULT 13 Backup, recovery, and replication 14 Data encryption 15 Securing information with Azure Key Vault 16 Azure Security Center and updates PART 4 - THE COOL STUFF 17 Machine learning and artificial intelligence 18 Azure Automation 19 Azure containers 20 Azure and the Internet of Things 21 Serverless computing

Machine Learning with Dynamics 365 and Power Platform - Aurelien Clere 2022-01-06

Apply cutting-edge AI techniques to your Dynamics 365 environment to create new solutions to old business problems In Machine Learning with Dynamics 365 and Power Platform: The Ultimate Guide to Apply Predictive Analytics, an accomplished team of digital and data analytics experts delivers a practical and comprehensive discussion of how to integrate AI Builder with Dataverse and Dynamics 365 to create real-world business solutions. It also walks you through how to build powerful machine learning models using Azure Data Lake, Databricks, Azure Synapse Analytics. The book is filled with clear explanations, visualizations, and working examples that get you up and running in your development of supervised, unsupervised, and reinforcement learning techniques using Microsoft machine learning tools and technologies. These strategies will transform your business verticals, reducing costs

and manual processes in finance and operations, retail, telecommunications, and manufacturing industries. The authors demonstrate: What machine learning is all about and how it can be applied to your organization's Dynamics 365 and Power Platform Projects The creation and management of environments for development, testing, and production of a machine learning project How adopting machine learning techniques will redefine the future of your ERP/CRM system Perfect for Technical Consultants, software developers, and solution architects, Machine Learning with Dynamics 365 and Power Platform is also an indispensable guide for Chief Technology Officers seeking an intuitive resource for how to implement machine learning in modern business applications to solve real-world problems.

[Azure Data Scientist Associate Certification Guide](#) - Andreas Botsikas 2021-12-03

Develop the skills you need to run machine learning workloads in Azure and pass the

DP-100 exam with ease Key Features Create end-to-end machine learning training pipelines, with or without code Track experiment progress using the cloud-based MLflow-compatible process of Azure ML services Operationalize your machine learning models by creating batch and real-time endpoints Book Description The Azure Data Scientist Associate Certification Guide helps you acquire practical knowledge for machine learning experimentation on Azure. It covers everything you need to pass the DP-100 exam and become a certified Azure Data Scientist Associate. Starting with an introduction to data science, you'll learn the terminology that will be used throughout the book and then move on to the Azure Machine Learning (Azure ML) workspace. You'll discover the studio interface and manage various components, such as data stores and compute clusters. Next, the book focuses on no-code and low-code experimentation, and shows you how to use the Automated ML wizard to locate and deploy

optimal models for your dataset. You'll also learn how to run end-to-end data science experiments using the designer provided in Azure ML Studio. You'll then explore the Azure ML Software Development Kit (SDK) for Python and advance to creating experiments and publishing models using code. The book also guides you in optimizing your model's hyperparameters using Hyperdrive before demonstrating how to use responsible AI tools to interpret and debug your models. Once you have a trained model, you'll learn to operationalize it for batch or real-time inferences and monitor it in production. By the end of this Azure certification study guide, you'll have gained the knowledge and the practical skills required to pass the DP-100 exam. What you will learn Create a working environment for data science workloads on Azure Run data experiments using Azure Machine Learning services Create training and inference pipelines using the designer or code Discover the best model for your dataset using Automated ML Use

hyperparameter tuning to optimize trained models
Deploy, use, and monitor models in production
Interpret the predictions of a trained model
Who this book is for This book is for developers who want to infuse their applications with AI capabilities and data scientists looking to scale their machine learning experiments in the Azure cloud. Basic knowledge of Python is needed to follow the code samples used in the book. Some experience in training machine learning models in Python using common frameworks like scikit-learn will help you understand the content more easily.

Mastering Azure Machine Learning Christoph Körner 2020-04-30

Master expert techniques for building automated and highly scalable end-to-end machine learning models and pipelines in Azure using TensorFlow, Spark, and Kubernetes
Key Features
Make sense of data on the cloud by implementing advanced analytics
Train and optimize advanced deep learning models efficiently on Spark using Azure

Databricks
Deploy machine learning models for batch and real-time scoring with Azure Kubernetes Service (AKS)
Book Description
The increase being seen in data volume today requires distributed systems, powerful algorithms, and scalable cloud infrastructure to compute insights and train and deploy machine learning (ML) models. This book will help you improve your knowledge of building ML models using Azure and end-to-end ML pipelines on the cloud. The book starts with an overview of an end-to-end ML project and a guide on how to choose the right Azure service for different ML tasks. It then focuses on Azure Machine Learning and takes you through the process of data experimentation, data preparation, and feature engineering using Azure Machine Learning and Python. You'll learn advanced feature extraction techniques using natural language processing (NLP), classical ML techniques, and the secrets of both a great recommendation engine and a performant

computer vision model using deep learning methods. You'll also explore how to train, optimize, and tune models using Azure Automated Machine Learning and HyperDrive, and perform distributed training on Azure. Then, you'll learn different deployment and monitoring techniques using Azure Kubernetes Services with Azure Machine Learning, along with the basics of MLOps—DevOps for ML to automate your ML process as CI/CD pipeline. By the end of this book, you'll have mastered Azure Machine Learning and be able to confidently design, build and operate scalable ML pipelines in Azure. What you will learn Setup your Azure Machine Learning workspace for data experimentation and visualization Perform ETL, data preparation, and feature extraction using Azure best practices Implement advanced feature extraction using NLP and word embeddings Train gradient boosted tree-ensembles, recommendation engines and deep neural networks on Azure Machine Learning Use

hyperparameter tuning and Azure Automated Machine Learning to optimize your ML models Employ distributed ML on GPU clusters using Horovod in Azure Machine Learning Deploy, operate and manage your ML models at scale Automated your end-to-end ML process as CI/CD pipelines for MLOps Who this book is for This machine learning book is for data professionals, data analysts, data engineers, data scientists, or machine learning developers who want to master scalable cloud-based machine learning architectures in Azure. This book will help you use advanced Azure services to build intelligent machine learning applications. A basic understanding of Python and working knowledge of machine learning are mandatory.

Hands-On Machine Learning with Microsoft Excel 2019 Julio Cesar Rodriguez Martino
2019-04-30

A practical guide to getting the most out of Excel, using it for data preparation, applying

machine learning models (including cloud services) and understanding the outcome of the data analysis. Key Features Use Microsoft's product Excel to build advanced forecasting models using varied examples Cover range of machine learning tasks such as data mining, data analytics, smart visualization, and more Derive data-driven techniques using Excel plugins and APIs without much code required Book Description We have made huge progress in teaching computers to perform difficult tasks, especially those that are repetitive and time-consuming for humans. Excel users, of all levels, can feel left behind by this innovation wave. The truth is that a large amount of the work needed to develop and use a machine learning model can be done in Excel. The book starts by giving a general introduction to machine learning, making every concept clear and understandable. Then, it shows every step of a machine learning project, from data collection, reading from different data sources, developing models, and

visualizing the results using Excel features and offerings. In every chapter, there are several examples and hands-on exercises that will show the reader how to combine Excel functions, add-ins, and connections to databases and to cloud services to reach the desired goal: building a full data analysis flow. Different machine learning models are shown, tailored to the type of data to be analyzed. At the end of the book, the reader is presented with some advanced use cases using Automated Machine Learning, and artificial neural network, which simplifies the analysis task and represents the future of machine learning. What you will learn Use Excel to preview and cleanse datasets Understand correlations between variables and optimize the input to machine learning models Use and evaluate different machine learning models from Excel Understand the use of different visualizations Learn the basic concepts and calculations to understand how artificial neural networks work Learn how to connect Excel to the

Microsoft Azure cloudGet beyond proof of concepts and build fully functional data analysis flowsWho this book is for This book is for data analysis, machine learning enthusiasts, project managers, and someone who doesn't want to code much for performing core tasks of machine learning. Each example will help you perform end-to-end smart analytics. Working knowledge of Excel is required.

Deep Learning in Natural Language Processing - Li Deng 2018-05-23

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.

Predictive Analytics with Microsoft Azure Machine Learning - Valentine Fontama

2014-11-25

Data Science and Machine Learning are in high demand, as customers are increasingly looking

for ways to glean insights from all their data. More customers now realize that Business Intelligence is not enough as the volume, speed and complexity of data now defy traditional analytics tools. While Business Intelligence addresses descriptive and diagnostic analysis, Data Science unlocks new opportunities through predictive and prescriptive analysis. The purpose of this book is to provide a gentle and instructionally organized introduction to the field of data science and machine learning, with a focus on building and deploying predictive models. The book also provides a thorough overview of the Microsoft Azure Machine Learning service using task oriented descriptions and concrete end-to-end examples, sufficient to ensure the reader can immediately begin using this important new service. It describes all aspects of the service from data ingress to applying machine learning and evaluating the resulting model, to deploying the resulting model as a machine learning web

service. Finally, this book attempts to have minimal dependencies, so that you can fairly easily pick and choose chapters to read. When dependencies do exist, they are listed at the start and end of the chapter. The simplicity of this new service from Microsoft will help to take Data Science and Machine Learning to a much broader audience than existing products in this space. Learn how you can quickly build and deploy sophisticated predictive models as machine learning web services with the new Azure Machine Learning service from Microsoft.

[Deep Learning on Windows](#) - Thimira Amaratunga 2021-02-25

Build deep learning and computer vision systems using Python, TensorFlow, Keras, OpenCV, and more, right within the familiar environment of Microsoft Windows. The book starts with an introduction to tools for deep learning and computer vision tasks followed by instructions to install, configure, and troubleshoot them. Here, you will learn how Python can help you build

deep learning models on Windows. Moving forward, you will build a deep learning model and understand the internal workings of a convolutional neural network on Windows. Further, you will go through different ways to visualize the internal workings of deep learning models along with an understanding of transfer learning where you will learn how to build a model architecture and use data augmentations. Next, you will manage and train deep learning models on Windows before deploying your application as a web application. You'll also do some basic image processing and work with computer vision options that will help you build various applications with deep learning. Finally, you will use generative adversarial networks along with reinforcement learning. After reading *Deep Learning on Windows*, you will be able to design deep learning models and web applications on the Windows operating system. What You Will Learn Get deep learning tools working on Microsoft Windows Understand

model visualization techniques, such as the built-in `plot_model` function of Keras and third-party visualization tools Build a robust training script Convert your deep learning model into a web application Generate handwritten digits with DCGAN (deep convolutional generative adversarial network) Understand the basics of reinforcement learning Who This Book Is For AI developers and enthusiasts wanting to work on the Windows platform.

Modern Computer Vision with PyTorch

Kishore Ayyadevara 2020-11-27

Get to grips with deep learning techniques for building image processing applications using PyTorch with the help of code notebooks and test questions Key Features Implement solutions to 50 real-world computer vision applications using PyTorch Understand the theory and working mechanisms of neural network architectures and their implementation Discover best practices using a custom library created especially for this book Book Description Deep

learning is the driving force behind many recent advances in various computer vision (CV) applications. This book takes a hands-on approach to help you to solve over 50 CV problems using PyTorch1.x on real-world datasets. You'll start by building a neural network (NN) from scratch using NumPy and PyTorch and discover best practices for tweaking its hyperparameters. You'll then perform image classification using convolutional neural networks and transfer learning and understand how they work. As you progress, you'll implement multiple use cases of 2D and 3D multi-object detection, segmentation, human-pose-estimation by learning about the R-CNN family, SSD, YOLO, U-Net architectures, and the Detectron2 platform. The book will also guide you in performing facial expression swapping, generating new faces, and manipulating facial expressions as you explore autoencoders and modern generative adversarial networks. You'll learn how to combine CV with NLP techniques,

such as LSTM and transformer, and RL techniques, such as Deep Q-learning, to implement OCR, image captioning, object detection, and a self-driving car agent. Finally, you'll move your NN model to production on the AWS Cloud. By the end of this book, you'll be able to leverage modern NN architectures to solve over 50 real-world CV problems confidently. What you will learn

- Train a NN from scratch with NumPy and PyTorch
- Implement 2D and 3D multi-object detection and segmentation
- Generate digits and DeepFakes with autoencoders and advanced GANs
- Manipulate images using CycleGAN, Pix2PixGAN, StyleGAN2, and SRGAN
- Combine CV with NLP to perform OCR, image captioning, and object detection
- Combine CV with reinforcement learning to build agents that play pong and self-drive a car
- Deploy a deep learning model on the AWS server using FastAPI and Docker
- Implement over 35 NN architectures and common OpenCV utilities

Who this book is for

This book is for beginners to PyTorch and intermediate-level machine learning practitioners who are looking to get well-versed with computer vision techniques using deep learning and PyTorch. If you are just getting started with neural networks, you'll find the use cases accompanied by notebooks in GitHub present in this book useful. Basic knowledge of the Python programming language and machine learning is all you need to get started with this book.

Engineering MLOps - Emmanuel Raj

2021-04-19

Get up and running with machine learning life cycle management and implement MLOps in your organization
Key Features
Become well-versed with MLOps techniques to monitor the quality of machine learning models in production
Explore a monitoring framework for ML models in production and learn about end-to-end traceability for deployed models
Perform CI/CD to automate new implementations in ML

pipelines
Book Description
Engineering MLps presents comprehensive insights into MLOps coupled with real-world examples in Azure to help you to write programs, train robust and scalable ML models, and build ML pipelines to train and deploy models securely in production. The book begins by familiarizing you with the MLOps workflow so you can start writing programs to train ML models. Then you'll then move on to explore options for serializing and packaging ML models post-training to deploy them to facilitate machine learning inference, model interoperability, and end-to-end model traceability. You'll learn how to build ML pipelines, continuous integration and continuous delivery (CI/CD) pipelines, and monitor pipelines to systematically build, deploy, monitor, and govern ML solutions for businesses and industries. Finally, you'll apply the knowledge you've gained to build real-world projects. By the end of this ML book, you'll have a 360-degree view of MLOps and be ready to implement

MLOps in your organization. What you will learn
Formulate data governance strategies and pipelines for ML training and deployment
Get to grips with implementing ML pipelines, CI/CD pipelines, and ML monitoring pipelines
Design a robust and scalable microservice and API for test and production environments
Curate your custom CD processes for related use cases and organizations
Monitor ML models, including monitoring data drift, model drift, and application performance
Build and maintain automated ML systems
Who this book is for
This MLOps book is for data scientists, software engineers, DevOps engineers, machine learning engineers, and business and technology leaders who want to build, deploy, and maintain ML systems in production using MLOps principles and techniques. Basic knowledge of machine learning is necessary to get started with this book.

Cognitive Computing Recipes - Adnan Masood
2019-03-27

Solve your AI and machine learning problems using complete and real-world code examples. Using a problem-solution approach, this book makes deep learning and machine learning accessible to everyday developers, by providing a combination of tools such as cognitive services APIs, machine learning platforms, and libraries. Along with an overview of the contemporary technology landscape, *Machine Learning and Deep Learning with Cognitive Computing Recipes* covers the business case for machine learning and deep learning. Covering topics such as digital assistants, computer vision, text analytics, speech, and robotics process automation this book offers a comprehensive toolkit that you can apply quickly and easily in your own projects. With its focus on Microsoft Cognitive Services offerings, you'll see recipes using multiple different environments including TensorFlow and CNTK to give you a broader perspective of the deep learning ecosystem. What You Will Learn
Build production-ready

solutions using Microsoft Cognitive Services APIs Apply deep learning using TensorFlow and Microsoft Cognitive Toolkit (CNTK) Solve enterprise problems in natural language processing and computer vision Discover the machine learning development life cycle - from formal problem definition to deployment at scale Who This Book Is For Software engineers and enterprise architects who wish to understand machine learning and deep learning by building applications and solving real-world business problems.

Deep Learning with Azure - Mathew Salvaris
2018-08-24

Get up-to-speed with Microsoft's AI Platform. Learn to innovate and accelerate with open and powerful tools and services that bring artificial intelligence to every data scientist and developer. Artificial Intelligence (AI) is the new normal. Innovations in deep learning algorithms and hardware are happening at a rapid pace. It is no longer a question of should I build AI into

my business, but more about where do I begin and how do I get started with AI? Written by expert data scientists at Microsoft, Deep Learning with the Microsoft AI Platform helps you with the how-to of doing deep learning on Azure and leveraging deep learning to create innovative and intelligent solutions. Benefit from guidance on where to begin your AI adventure, and learn how the cloud provides you with all the tools, infrastructure, and services you need to do AI. What You'll Learn Become familiar with the tools, infrastructure, and services available for deep learning on Microsoft Azure such as Azure Machine Learning services and Batch AI Use pre-built AI capabilities (Computer Vision, OCR, gender, emotion, landmark detection, and more) Understand the common deep learning models, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), generative adversarial networks (GANs) with sample code and understand how the field is evolving Discover the options for training and

operationalizing deep learning models on Azure
Who This Book Is For Professional data scientists who are interested in learning more about deep learning and how to use the Microsoft AI platform. Some experience with Python is helpful.

Automated Machine Learning with Microsoft Azure - Dennis Michael Sawyers
2021-04-23

A practical, step-by-step guide to using Microsoft's AutoML technology on the Azure Machine Learning service for developers and data scientists working with the Python programming language
Key Features: Create, deploy, productionalize, and scale automated machine learning solutions on Microsoft Azure
Improve the accuracy of your ML models through automatic data featurization and model training
Increase productivity in your organization by using artificial intelligence to solve common problems
Book Description: Automated Machine Learning with Microsoft

Azure helps you to build high-performing, accurate machine learning models in record time. It allows anyone to easily harness the power of artificial intelligence and increase the productivity and profitability of your business. With a series of clicks on a guided user interface (GUI), novices and seasoned data scientists alike can train and deploy machine learning solutions to production with ease. This book will teach you how to use Azure AutoML with both the GUI as well as the AzureML Python software development kit (SDK) in a careful, step-by-step way. First, you'll learn how to prepare data, train models, and register them to your Azure Machine Learning workspace. You'll then discover how to take those models and use them to create both automated batch solutions using machine learning pipelines and real-time scoring solutions using Azure Kubernetes Service (AKS). Finally, you will be able to use AutoML on your own data to not only train regression, classification, and forecasting models but also

use them to solve a wide variety of business problems. By the end of this Azure book, you'll be able to show your business partners exactly how your ML models are making predictions through automatically generated charts and graphs, earning their trust and respect. What You Will Learn: Understand how to train classification, regression, and forecasting ML algorithms with Azure AutoML Prepare data for Azure AutoML to ensure smooth model training and deployment Adjust AutoML configuration settings to make your models as accurate as possible Determine when to use a batch-scoring solution versus a real-time scoring solution Productionalize your AutoML solution with Azure Machine Learning pipelines Create real-time scoring solutions with AutoML and Azure Kubernetes Service Discover how to quickly deliver value and earn business trust using AutoML Train a large number of AutoML models at once using the AzureML Python SDK Who this book is for: Data scientists, aspiring data

scientists, machine learning engineers, or anyone interested in applying artificial intelligence or machine learning in their business will find this book useful. You need to have beginner-level knowledge of artificial intelligence and a technical background in computer science, statistics, or information technology before getting started with this machine learning book. Familiarity with Python will help you implement this book's more advanced features, but even data analysts and SQL experts will be able to train ML models after finishing this book.

Deep Learning with Microsoft Cognitive Toolkit Quick Start Gui deWillem Meints 2019-03-28

Learn how to train popular deep learning architectures such as autoencoders, convolutional and recurrent neural networks while discovering how you can use deep learning models in your software applications with Microsoft Cognitive Toolkit Key FeaturesUnderstand the fundamentals of

Microsoft Cognitive Toolkit and set up the development environment Train different types of neural networks using Cognitive Toolkit and deploy it to production Evaluate the performance of your models and improve your deep learning skills

Book Description Cognitive Toolkit is a very popular and recently open sourced deep learning toolkit by Microsoft. Cognitive Toolkit is used to train fast and effective deep learning models. This book will be a quick introduction to using Cognitive Toolkit and will teach you how to train and validate different types of neural networks, such as convolutional and recurrent neural networks. This book will help you understand the basics of deep learning. You will learn how to use Microsoft Cognitive Toolkit to build deep learning models and discover what makes this framework unique so that you know when to use it. This book will be a quick, no-nonsense introduction to the library and will teach you how to train different types of neural networks, such as convolutional neural

networks, recurrent neural networks, autoencoders, and more, using Cognitive Toolkit. Then we will look at two scenarios in which deep learning can be used to enhance human capabilities. The book will also demonstrate how to evaluate your models' performance to ensure it trains and runs smoothly and gives you the most accurate results. Finally, you will get a short overview of how Cognitive Toolkit fits in to a DevOps environment

What you will learn Set up your deep learning environment for the Cognitive Toolkit on Windows and Linux Pre-process and feed your data into neural networks Use neural networks to make efficient predictions and recommendations Train and deploy efficient neural networks such as CNN and RNN Detect problems in your neural network using TensorBoard Integrate Cognitive Toolkit with Azure ML Services for effective deep learning

Who this book is for Data Scientists, Machine learning developers, AI developers who wish to train and deploy

effective deep learning models using Microsoft CNTK will find this book to be useful. Readers

need to have experience in Python or similar object-oriented language like C# or Java.