

Real World Machine Learning

Turning text into valuable information is essential for businesses looking to gain a competitive advantage. With recent improvements in natural language processing (NLP), users now have many options for solving complex challenges. But it's not always clear which NLP tools or libraries would work for a business's needs, or which techniques you should use and in what order. This practical book provides data scientists and developers with blueprints for best practice solutions to common tasks in text analytics and natural language processing. Authors Jens Albrecht, Sidharth Ramachandran, and Christian Winkler provide real-world case studies and detailed code examples in Python to help you get started quickly. Extract data from APIs and web pages Prepare textual data for statistical analysis and machine learning Use machine learning for classification, topic modeling, and summarization Explain AI models and classification results Explore and visualize semantic similarities with word embeddings Identify customer sentiment in product reviews Create a knowledge graph based on named entities and their relations 100 recipes that teach you how to perform various machine learning tasks in the real world About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide Learn about perceptrons and see how they are used to build neural networks Stuck while making sense of images, text, speech, and real estate? This guide will come to your rescue, showing you how to perform machine learning for each one of these using various techniques Who This Book Is For This book is for Python programmers who are looking to use machine-learning algorithms to create real-world applications. This book is friendly to Python beginners, but familiarity with Python programming would certainly be useful to play around with the code. What You Will Learn Explore classification algorithms and apply them to the income bracket estimation problem Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Explore data visualization techniques to interact with your data in diverse ways Find out how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Analyze stock market data using Conditional Random Fields Work with image data and build systems for image recognition and biometric face recognition Grasp how to use deep neural networks to build an optical character recognition system In Detail Machine learning is becoming increasingly pervasive in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. With this book, you will learn how to perform various machine learning tasks in different environments. We'll start by exploring a range of real-life scenarios where machine learning can be used, and look at various building blocks. Throughout the book, you'll use a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. You'll discover how to deal with various types of data and explore the differences between machine learning paradigms such as supervised and unsupervised learning. We also cover a range of regression techniques, classification algorithms, predictive modeling, data visualization techniques, recommendation engines, and more with the help of real-world examples. Style and approach You will explore various real-life scenarios in this book where machine learning can be used, and learn about different building blocks of machine learning using independent recipes in the book.

This book includes 9 projects on building smart and practical AI-based systems. These projects cover solutions to different domain-specific problems in healthcare, e-commerce and more. With this book, you will apply different machine learning and deep learning techniques and learn how to build your own intelligent applications for smart ...

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to real-world problems Explore data visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

"Machine learning algorithms and research are mushrooming due to their accuracy at solving problems. This course walks you through developing real-world projects using TensorFlow in your ML projects. The initial project will deal with assessing the viability of expanding your Restaurant business using a single variable linear regression. You will use Linear Regression with multiple variables with an example involving buying and selling a property at the best prices and use a dataset containing 11 features to deal with it. Next, you will create an algorithm to detect anomalous behavior in server computers using Gaussian methods. Finally, you'll design and build a convolutional Neural Networks model on a Traffic Signal Classifier from scratch. By the end of this course you will be using TensorFlow in real-world scenarios, and you'll be confident enough to use ML Algorithms to build your own projects."--Resource description page.

Powerful smart applications using deep learning algorithms to dominate numerical computing, deep learning, and functional programming. Key Features Explore machine learning techniques with prominent open source Scala libraries such as Spark ML, H2O, MXNet, Zeppelin, and DeepLearning4j Solve real-world machine learning problems by delving complex numerical

computing with Scala functional programming in a scalable and faster way Cover all key aspects such as collection, storing, processing, analyzing, and evaluation required to build and deploy machine models on computing clusters using Scala Play framework. Book Description Machine learning has had a huge impact on academia and industry by turning data into actionable information. Scala has seen a steady rise in adoption over the past few years, especially in the fields of data science and analytics. This book is for data scientists, data engineers, and deep learning enthusiasts who have a background in complex numerical computing and want to know more hands-on machine learning application development. If you're well versed in machine learning concepts and want to expand your knowledge by delving into the practical implementation of these concepts using the power of Scala, then this book is what you need! Through 11 end-to-end projects, you will be acquainted with popular machine learning libraries such as Spark ML, H2O, DeepLearning4j, and MXNet. At the end, you will be able to use numerical computing and functional programming to carry out complex numerical tasks to develop, build, and deploy research or commercial projects in a production-ready environment. What you will learn Apply advanced regression techniques to boost the performance of predictive models Use different classification algorithms for business analytics Generate trading strategies for Bitcoin and stock trading using ensemble techniques Train Deep Neural Networks (DNN) using H2O and Spark ML Utilize NLP to build scalable machine learning models Learn how to apply reinforcement learning algorithms such as Q-learning for developing ML application Learn how to use autoencoders to develop a fraud detection application Implement LSTM and CNN models using DeepLearning4j and MXNet Who this book is for If you want to leverage the power of both Scala and Spark to make sense of Big Data, then this book is for you. If you are well versed with machine learning concepts and wants to expand your knowledge by delving into the practical implementation using the power of Scala, then this book is what you need! Strong understanding of Scala Programming language is recommended. Basic familiarity with machine Learning techniques will be more helpful.

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Summary Machine Learning Systems: Designs that scale is an example-rich guide that teaches you how to implement reactive design solutions in your machine learning systems to make them as reliable as a well-built web app. Foreword by Sean Owen, Director of Data Science, Cloudera Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology If you're building machine learning models to be used on a small scale, you don't need this book. But if you're a developer building a production-grade ML application that needs quick response times, reliability, and good user experience, this is the book for you. It collects principles and practices of machine learning systems that are dramatically easier to run and maintain, and that are reliably better for users. About the Book Machine Learning Systems: Designs that scale teaches you to design and implement production-ready ML systems. You'll learn the principles of reactive design as you build pipelines with Spark, create highly scalable services with Akka, and use powerful machine learning libraries like MLlib on massive datasets. The examples use the Scala language, but the same ideas and tools work in Java, as well. What's Inside Working with Spark, MLlib, and Akka Reactive design patterns Monitoring and maintaining a large-scale system Futures, actors, and supervision About the Reader Readers need intermediate skills in Java or Scala. No prior machine learning experience is assumed. About the Author Jeff Smith builds powerful machine learning systems. For the past decade, he has been working on building data science applications, teams, and companies as part of various teams in New York, San Francisco, and Hong Kong. He blogs (<https://medium.com/@jeffksmithjr>), tweets (@jeffksmithjr), and speaks (www.jeffsmith.tech/speaking) about various aspects of building real-world machine learning systems. Table of Contents PART 1 - FUNDAMENTALS OF REACTIVE MACHINE LEARNING Learning reactive machine learning Using reactive tools PART 2 - BUILDING A REACTIVE MACHINE LEARNING SYSTEM Collecting data Generating features Learning models Evaluating models Publishing models Responding PART 3 - OPERATING A MACHINE LEARNING SYSTEM Delivering Evolving intelligence

There is a significant deficiency among contemporary medicine practices reflected by experts making medical decisions for a large proportion of the population for which no or minimal data exists. Fortunately, our capacity to procure and apply such information is rapidly rising. As medicine becomes more individualized, the implementation of health IT and data interoperability become essential components to delivering quality healthcare. Quality Assurance in the Era of Individualized Medicine is a collection of innovative research on the methods and utilization of digital readouts to fashion an individualized therapy instead of a mass-population-directed strategy. While highlighting topics including assistive technologies, patient management, and clinical practices, this book is ideally designed for health professionals, doctors, nurses, hospital management, medical administrators, IT specialists, data scientists, researchers, academicians, and

students.

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.

This book provides an overview of Federated Learning and how it can be used to build real-world AI-enabled applications. Real-world AI applications frequently have training data distributed in many different locations, with data at different sites having different properties and different formats. In many cases, data movement is not permitted due to security concerns, bandwidth, cost or regulatory restriction. Under these conditions, techniques of federated learning can enable creation of practical applications. Creating practical applications requires implementation of the cycle of learning from data, inferring from data, and acting based on the inference. This book will be the first one to cover all stages of the Learn-Infer-Act cycle, and presents a set of patterns to apply federation to all stages. Another distinct feature of the book is the use of real-world applications with an approach that discusses all aspects that need to be considered in an operational system, including handling of data issues during federation, maintaining compliance with enterprise security policies, and simplifying the logistics of federated AI in enterprise contexts. The book considers federation from a manner agnostic to the actual AI models, allowing the concepts to be applied to all varieties of AI models. This book is probably the first one to cover the space of enterprise AI-based applications in a holistic manner.

Artificial intelligence and its various components are rapidly engulfing almost every professional industry. Specific features of AI that have proven to be vital solutions to numerous real-world issues are machine learning and deep learning. These intelligent agents unlock higher levels of performance and efficiency, creating a wide span of industrial applications. However, there is a lack of research on the specific uses of machine/deep learning in the professional realm. Machine Learning and Deep Learning in Real-Time Applications provides emerging research exploring the theoretical and practical aspects of machine learning and deep learning and their implementations as well as their ability to solve real-world problems within several professional disciplines including healthcare, business, and computer science. Featuring coverage on a broad range of topics such as image processing, medical improvements, and smart grids, this book is ideally designed for researchers, academicians, scientists, industry experts, scholars, IT professionals, engineers, and students seeking current research on the multifaceted uses and implementations of machine learning and deep learning across the globe.

How can you successfully deploy AI? When AI works, it's nothing short of brilliant, helping companies make or save tremendous amounts of money while delighting customers on an unprecedented scale. When it fails, the results can be devastating. Most AI models never make it out of testing, but those failures aren't random. This practical guide to deploying AI lays out a human-first, responsible approach that has seen more than three times the success rate when compared to the industry average. In Real World AI, Alyssa Simpson Rochwerger and Wilson Pang share dozens of AI stories from startups and global enterprises alike featuring personal experiences from people who have worked on global AI deployments that impact billions of people every day. ? AI for business doesn't have to be overwhelming. Real World AI uses plain language to walk you through an AI approach that you can feel confident about-for your business and for your customers.

Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. In Real-world Natural Language Processing you will learn how to: Design, develop, and deploy useful NLP applications Create named entity taggers Build machine translation systems Construct language generation systems and chatbots Use advanced NLP concepts such as attention and transfer learning Real-world Natural Language Processing teaches you how to create practical NLP

applications without getting bogged down in complex language theory and the mathematics of deep learning. In this engaging book, you'll explore the core tools and techniques required to build a huge range of powerful NLP apps, including chatbots, language detectors, and text classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Training computers to interpret and generate speech and text is a monumental challenge, and the payoff for reducing labor and improving human/computer interaction is huge! The field of Natural Language Processing (NLP) is advancing rapidly, with countless new tools and practices. This unique book offers an innovative collection of NLP techniques with applications in machine translation, voice assistants, text generation, and more. About the book Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. Guided by clear explanations of each core NLP topic, you'll create many interesting applications including a sentiment analyzer and a chatbot. Along the way, you'll use Python and open source libraries like AllenNLP and HuggingFace Transformers to speed up your development process. What's inside Design, develop, and deploy useful NLP applications Create named entity taggers Build machine translation systems Construct language generation systems and chatbots About the reader For Python programmers. No prior machine learning knowledge assumed. About the author Masato Hagiwara received his computer science PhD from Nagoya University in 2009. He has interned at Google and Microsoft Research, and worked at Duolingo as a Senior Machine Learning Engineer. He now runs his own research and consulting company. Table of Contents PART 1 BASICS 1 Introduction to natural language processing 2 Your first NLP application 3 Word and document embeddings 4 Sentence classification 5 Sequential labeling and language modeling PART 2 ADVANCED MODELS 6 Sequence-to-sequence models 7 Convolutional neural networks 8 Attention and Transformer 9 Transfer learning with pretrained language models PART 3 PUTTING INTO PRODUCTION 10 Best practices in developing NLP applications 11 Deploying and serving NLP applications

Machine learning allows for non-conventional and productive answers for issues within various fields, including problems related to visually perceptive computers. Applying these strategies and algorithms to the area of computer vision allows for higher achievement in tasks such as spatial recognition, big data collection, and image processing. There is a need for research that seeks to understand the development and efficiency of current methods that enable machines to see. Challenges and Applications for Implementing Machine Learning in Computer Vision is a collection of innovative research that combines theory and practice on adopting the latest deep learning advancements for machines capable of visual processing. Highlighting a wide range of topics such as video segmentation, object recognition, and 3D modelling, this publication is ideally designed for computer scientists, medical professionals, computer engineers, information technology practitioners, industry experts, scholars, researchers, and students seeking current research on the utilization of evolving computer vision techniques.

Learn to solve challenging data science problems by building powerful machine learning models using Python About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide This practical tutorial tackles real-world computing problems through a rigorous and effective approach Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale Who This Book Is For This Learning Path is for Python programmers who are looking to use machine learning algorithms to create real-world applications. It is ideal for Python professionals who want to work with large and complex datasets and Python developers and analysts or data scientists who are looking to add to their existing skills by accessing some of the most powerful recent trends in data science. Experience with Python, Jupyter Notebooks, and command-line execution together with a good level of mathematical knowledge to understand the concepts is expected. Machine learning basic knowledge is also expected. What You Will Learn Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Apply your new-found skills to solve real problems, through clearly-explained code for every technique and test Compete with top data scientists by gaining a practical and theoretical understanding of cutting-edge deep learning algorithms Increase predictive accuracy with deep learning and scalable data-handling techniques Work with modern state-of-the-art large-scale machine learning techniques Learn to use Python code to implement a range of machine learning algorithms and techniques In Detail Machine learning is increasingly spreading in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. Machine learning is transforming the way we understand and interact with the world around us. In the first module, Python Machine Learning Cookbook, you will learn how to perform various machine learning tasks using a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. The second module, Advanced Machine Learning with Python, is designed to take you on a guided tour of the most relevant and powerful machine learning techniques and you'll acquire a broad set of powerful skills in the area of feature selection and feature engineering. The third module in this learning path, Large Scale Machine Learning with Python, dives into scalable machine learning and the three forms of scalability. It covers the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python. This Learning Path will teach you Python machine learning for the real world. The machine learning techniques covered in this Learning Path are at the forefront of commercial practice. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: Python Machine Learning Cookbook by Prateek Joshi Advanced Machine Learning with Python by John Hearty Large Scale Machine Learning with Python by Bastiaan Sjardin, Alberto Boschetti, Luca Massaron Style and approach This course is a smooth learning path that will teach you how to get started with Python machine learning for the real world, and develop solutions to real-world problems. Through this comprehensive course, you'll learn to create the most effective machine learning techniques from scratch and more!

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can

focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow

Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key Features Use machine learning and deep learning principles to build real-world projects Get to grips with TensorFlow's impressive range of module offerings Implement projects on GANs, reinforcement learning, and capsule network Book Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits—simplicity, efficiency, and flexibility—of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learn Understand the TensorFlow ecosystem using various datasets and techniques Create recommendation systems for quality product recommendations Build projects using CNNs, NLP, and Bayesian neural networks Play Pac-Man using deep reinforcement learning Deploy scalable TensorFlow-based machine learning systems Generate your own book script using RNNs Who this book is for TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In

The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

Real-World Machine Learning

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

Tackle the real-world complexities of modern machine learning with innovative, cutting-edge, techniques About This Book Fully-coded working examples using a wide range of machine learning libraries and tools, including Python, R, Julia, and Spark Comprehensive practical solutions taking you into the future of machine learning Go a step further and integrate your machine learning projects with Hadoop Who This Book Is For This book has been created for data scientists who want to see machine learning in action and explore its real-world application. With guidance on everything from the fundamentals of machine learning and predictive analytics to the latest innovations set to lead the big data revolution into the future, this is an unmissable resource for anyone dedicated to tackling current big data challenges. Knowledge of programming (Python and R) and mathematics is advisable if you want to get started immediately. What You Will Learn Implement a wide range of algorithms and techniques for tackling complex data Get to grips with some of the most powerful languages in data science, including R, Python, and Julia Harness the capabilities of Spark and Hadoop to manage and process data

successfully Apply the appropriate machine learning technique to address real-world problems Get acquainted with Deep learning and find out how neural networks are being used at the cutting-edge of machine learning Explore the future of machine learning and dive deeper into polyglot persistence, semantic data, and more In Detail Finding meaning in increasingly larger and more complex datasets is a growing demand of the modern world. Machine learning and predictive analytics have become the most important approaches to uncover data gold mines. Machine learning uses complex algorithms to make improved predictions of outcomes based on historical patterns and the behaviour of data sets. Machine learning can deliver dynamic insights into trends, patterns, and relationships within data, immensely valuable to business growth and development. This book explores an extensive range of machine learning techniques uncovering hidden tricks and tips for several types of data using practical and real-world examples. While machine learning can be highly theoretical, this book offers a refreshing hands-on approach without losing sight of the underlying principles. Inside, a full exploration of the various algorithms gives you high-quality guidance so you can begin to see just how effective machine learning is at tackling contemporary challenges of big data. This is the only book you need to implement a whole suite of open source tools, frameworks, and languages in machine learning. We will cover the leading data science languages, Python and R, and the underrated but powerful Julia, as well as a range of other big data platforms including Spark, Hadoop, and Mahout. Practical Machine Learning is an essential resource for the modern data scientists who want to get to grips with its real-world application. With this book, you will not only learn the fundamentals of machine learning but dive deep into the complexities of real world data before moving on to using Hadoop and its wider ecosystem of tools to process and manage your structured and unstructured data. You will explore different machine learning techniques for both supervised and unsupervised learning; from decision trees to Naive Bayes classifiers and linear and clustering methods, you will learn strategies for a truly advanced approach to the statistical analysis of data. The book also explores the cutting-edge advancements in machine learning, with worked examples and guidance on deep learning and reinforcement learning, providing you with practical demonstrations and samples that help take the theory—and mystery—out of even the most advanced machine learning methodologies. Style and approach A practical data science tutorial designed to give you an insight into the practical application of machine learning, this book takes you through complex concepts and tasks in an accessible way. Featuring information on a wide range of data science techniques, Practical Machine Learning is a comprehensive data science resource.

In a world where big data is the norm and near-real-time decisions are crucial, machine learning (ML) is a critical component of the data workflow. Machine learning systems can quickly crunch massive amounts of information to offer insights and make decisions in a way that matches or even surpasses human cognitive abilities. These systems use sophisticated computational and statistical tools to build models that can recognize and visualize patterns, predict outcomes, forecast values, and make recommendations. Real-World Machine Learning is a practical guide designed to teach developers the art of ML project execution. The book introduces the day-to-day practice of machine learning and prepares readers to successfully build and deploy powerful ML systems. Using the Python language and the R statistical package, it starts with core concepts like data acquisition and modeling, classification, and regression. Then it moves through the most important ML tasks, like model validation, optimization and feature engineering. It uses real-world examples that help readers anticipate and overcome common pitfalls. Along the way, they will discover scalable and online algorithms for large and streaming data sets. Advanced readers will appreciate the in-depth discussion of enhanced ML systems through advanced data exploration and pre-processing methods. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

This hands-on book will help you make your machine learning models fairer, safer, and more reliable and in turn improve business outcomes. Every chapter introduces a new mission where you learn how to apply interpretation methods to realistic use cases with methods that work for any model type as well as methods specific for deep neural networks.

Learn the skills necessary to design, build, and deploy applications powered by machine learning (ML). Through the course of this hands-on book, you'll build an example ML-driven application from initial idea to deployed product. Data scientists, software engineers, and product managers—including experienced practitioners and novices alike—will learn the tools, best practices, and challenges involved in building a real-world ML application step by step. Author Emmanuel Ameisen, an experienced data scientist who led an AI education program, demonstrates practical ML concepts using code snippets, illustrations, screenshots, and interviews with industry leaders. Part I teaches you how to plan an ML application and measure success. Part II explains how to build a working ML model. Part III demonstrates ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you: Define your product goal and set up a machine learning problem Build your first end-to-end pipeline quickly and acquire an initial dataset Train and evaluate your ML models and address performance bottlenecks Deploy and monitor your models in a production environment

Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks. Intended for those who want to learn how to use R's machine learning capabilities and gain insight from your data. Perhaps you already know a bit about machine learning, but have never used R; or perhaps you know a little R but are new to machine learning. In either case, this book will get you up and running quickly. It would be helpful to have a bit of familiarity with basic programming concepts, but no prior experience is required.

Lifelong Machine Learning, Second Edition is an introduction to an advanced machine learning paradigm that continuously learns by accumulating past knowledge that it then uses in future learning and problem solving. In contrast, the current dominant machine learning paradigm learns in isolation: given a training dataset, it runs a machine learning algorithm on the dataset to produce a model that is then used in its intended application. It makes no attempt to retain the learned knowledge and use it in subsequent learning. Unlike this isolated system, humans learn effectively with only a few examples precisely because our learning is very knowledge-driven: the knowledge learned in the past helps us learn new things with little data or effort. Lifelong learning aims to emulate this capability, because without it, an AI system cannot be considered truly intelligent. Research in lifelong learning has developed significantly in the relatively short time since the first edition of this book was published. The purpose of this second edition is to expand the definition of lifelong learning, update the content of several chapters, and add a new chapter about continual learning in deep neural networks—which has been actively researched over the past two or three years. A few chapters have also been reorganized to make each of them more coherent for the reader. Moreover, the authors want to propose a unified framework for the research area.

Currently, there are several research topics in machine learning that are closely related to lifelong learning—most notably, multi-task learning, transfer learning, and meta-learning—because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. Its goal is to introduce this emerging machine learning paradigm and present a comprehensive survey and review of the important research results and latest ideas in the area. This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book for courses in any of these related fields.

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

Create AI applications in Python and lay the foundations for your career in data science Key Features Practical examples that explain key machine learning algorithms Explore neural networks in detail with interesting examples Master core AI concepts with engaging activities Book Description Machine learning and neural networks are pillars on which you can build intelligent applications. *Artificial Intelligence and Machine Learning Fundamentals* begins by introducing you to Python and discussing AI search algorithms. You will cover in-depth mathematical topics, such as regression and classification, illustrated by Python examples. As you make your way through the book, you will progress to advanced AI techniques and concepts, and work on real-life datasets to form decision trees and clusters. You will be introduced to neural networks, a powerful tool based on Moore's law. By the end of this book, you will be confident when it comes to building your own AI applications with your newly acquired skills! What you will learn Understand the importance, principles, and fields of AI Implement basic artificial intelligence concepts with Python Apply regression and classification concepts to real-world problems Perform predictive analysis using decision trees and random forests Carry out clustering using the k-means and mean shift algorithms Understand the fundamentals of deep learning via practical examples Who this book is for *Artificial Intelligence and Machine Learning Fundamentals* is for software developers and data scientists who want to enrich their projects with machine learning. You do not need any prior experience in AI. However, it's recommended that you have knowledge of high school-level mathematics and at least one programming language (preferably Python).

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

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

With the flexibility and features of scikit-learn and Python, build machine learning algorithms that optimize the programming process and take application performance to a whole new level Key Features Explore scikit-learn uniform API and its application into any type of model Understand the difference between supervised and unsupervised models Learn the usage of machine learning through real-world examples Book Description As machine learning algorithms become popular, new tools that optimize these algorithms are also developed. *Machine Learning Fundamentals* explains you how to use the syntax of scikit-learn. You'll study the difference between supervised and unsupervised models, as well as the importance of choosing the appropriate algorithm for each dataset. You'll apply unsupervised clustering algorithms over real-world datasets, to discover patterns and profiles, and explore the process to solve an unsupervised machine learning problem. The focus of the book then shifts to supervised learning algorithms. You'll learn to implement different supervised algorithms and develop neural network structures using the scikit-learn package. You'll also learn how to perform coherent result analysis to improve the performance of the algorithm by tuning hyperparameters. By the end of this book, you will have gain all the skills required to start programming machine learning algorithms. What you will learn Understand the importance of data representation Gain insights into the differences between supervised and unsupervised models Explore data using the Matplotlib library Study popular algorithms, such as k-means, Mean-Shift, and DBSCAN Measure model performance through different metrics Implement a confusion matrix using scikit-learn Study popular algorithms, such as Naïve-Bayes, Decision Tree, and SVM Perform error analysis to

improve the performance of the model Learn to build a comprehensive machine learning program Who this book is for Machine Learning Fundamentals is designed for developers who are new to the field of machine learning and want to learn how to use the scikit-learn library to develop machine learning algorithms. You must have some knowledge and experience in Python programming, but you do not need any prior knowledge of scikit-learn or machine learning algorithms.

Understand the fundamentals of machine learning with R and build your own dynamic algorithms to tackle complicated real-world problems successfully About This Book Get to grips with the concepts of machine learning through exciting real-world examples Visualize and solve complex problems by using power-packed R constructs and its robust packages for machine learning Learn to build your own machine learning system with this example-based practical guide Who This Book Is For If you are interested in mining useful information from data using state-of-the-art techniques to make data-driven decisions, this is a go-to guide for you. No prior experience with data science is required, although basic knowledge of R is highly desirable. Prior knowledge in machine learning would be helpful but is not necessary. What You Will Learn Utilize the power of R to handle data extraction, manipulation, and exploration techniques Use R to visualize data spread across multiple dimensions and extract useful features Explore the underlying mathematical and logical concepts that drive machine learning algorithms Dive deep into the world of analytics to predict situations correctly Implement R machine learning algorithms from scratch and be amazed to see the algorithms in action Write reusable code and build complete machine learning systems from the ground up Solve interesting real-world problems using machine learning and R as the journey unfolds Harness the power of robust and optimized R packages to work on projects that solve real-world problems in machine learning and data science In Detail Data science and machine learning are some of the top buzzwords in the technical world today. From retail stores to Fortune 500 companies, everyone is working hard to making machine learning give them data-driven insights to grow their business. With powerful data manipulation features, machine learning packages, and an active developer community, R empowers users to build sophisticated machine learning systems to solve real-world data problems. This book takes you on a data-driven journey that starts with the very basics of R and machine learning and gradually builds upon the concepts to work on projects that tackle real-world problems. You'll begin by getting an understanding of the core concepts and definitions required to appreciate machine learning algorithms and concepts. Building upon the basics, you will then work on three different projects to apply the concepts of machine learning, following current trends and cover major algorithms as well as popular R packages in detail. These projects have been neatly divided into six different chapters covering the worlds of e-commerce, finance, and social-media, which are at the very core of this data-driven revolution. Each of the projects will help you to understand, explore, visualize, and derive insights depending upon the domain and algorithms. Through this book, you will learn to apply the concepts of machine learning to deal with data-related problems and solve them using the powerful yet simple language, R. Style and approach The book is an enticing journey that starts from the very basics to gradually pick up pace as the story unfolds. Each concept is first defined in the larger context of things succinctly, followed by a detailed explanation of their application. Each topic is explained with the help of a project that solves a real real-world problem involving hands-on work thus giving you a deep insight into the world of machine learning.

Time to flex your machine learning muscles! Take on the carefully designed challenges of the Machine Learning Bookcamp and master essential ML techniques through practical application. Summary In Machine Learning Bookcamp you will: Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Apply ML to complex datasets with images Deploy ML models to a production-ready environment The only way to learn is to practice! In Machine Learning Bookcamp, you'll create and deploy Python-based machine learning models for a variety of increasingly challenging projects. Taking you from the basics of machine learning to complex applications such as image analysis, each new project builds on what you've learned in previous chapters. You'll build a portfolio of business-relevant machine learning projects that hiring managers will be excited to see. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Master key machine learning concepts as you build actual projects! Machine learning is what you need for analyzing customer behavior, predicting price trends, evaluating risk, and much more. To master ML, you need great examples, clear explanations, and lots of practice. This book delivers all three! About the book Machine Learning Bookcamp presents realistic, practical machine learning scenarios, along with crystal-clear coverage of key concepts. In it, you'll complete engaging projects, such as creating a car price predictor using linear regression and deploying a churn prediction service. You'll go beyond the algorithms and explore important techniques like deploying ML applications on serverless systems and serving models with Kubernetes and Kubeflow. Dig in, get your hands dirty, and have fun building your ML skills! What's inside Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Deploy ML models to a production-ready environment About the reader Python programming skills assumed. No previous machine learning knowledge is required. About the author Alexey Grigorev is a principal data scientist at OLX Group. He runs DataTalks.Club, a community of people who love data. Table of Contents 1 Introduction to machine learning 2 Machine learning for regression 3 Machine learning for classification 4 Evaluation metrics for classification 5 Deploying machine learning models 6 Decision trees and ensemble learning 7 Neural networks and deep learning 8 Serverless deep learning 9 Serving models with Kubernetes and Kubeflow

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

In recent years, machine learning has gained a lot of interest. Due to the advances in processor technology and the availability of large amounts of data, machine learning techniques have provided astounding results in areas such as object recognition or natural language processing. New approaches, e.g. deep learning, have provided groundbreaking outcomes in fields such as

multimedia mining or voice recognition. Machine learning is now used in virtually every domain and deep learning algorithms are present in many devices such as smartphones, cars, drones, healthcare equipment, or smart home devices. The Internet, cloud computing and the Internet of Things produce a tsunami of data and machine learning provides the methods to effectively analyze the data and discover actionable knowledge. This book describes the most common machine learning techniques such as Bayesian models, support vector machines, decision tree induction, regression analysis, and recurrent and convolutional neural networks. It first gives an introduction into the principles of machine learning. It then covers the basic methods including the mathematical foundations. The biggest part of the book provides common machine learning algorithms and their applications. Finally, the book gives an outlook into some of the future developments and possible new research areas of machine learning and artificial intelligence in general. This book is meant to be an introduction into machine learning. It does not require prior knowledge in this area. It covers some of the basic mathematical principle but intends to be understandable even without a background in mathematics. It can be read chapter wise and intends to be comprehensible, even when not starting in the beginning. Finally, it also intends to be a reference book. Key Features: Describes real world problems that can be solved using Machine Learning Provides methods for directly applying Machine Learning techniques to concrete real world problems Demonstrates how to apply Machine Learning techniques using different frameworks such as TensorFlow, MALLETT, R

Unleash the power of Python and its robust data science capabilities About This Book Unleash the power of Python 3 objects Learn to use powerful Python libraries for effective data processing and analysis Harness the power of Python to analyze data and create insightful predictive models Unlock deeper insights into machine learning with this vital guide to cutting-edge predictive analytics Who This Book Is For Entry-level analysts who want to enter in the data science world will find this course very useful to get themselves acquainted with Python's data science capabilities for doing real-world data analysis. What You Will Learn Install and setup Python Implement objects in Python by creating classes and defining methods Get acquainted with NumPy to use it with arrays and array-oriented computing in data analysis Create effective visualizations for presenting your data using Matplotlib Process and analyze data using the time series capabilities of pandas Interact with different kind of database systems, such as file, disk format, Mongo, and Redis Apply data mining concepts to real-world problems Compute on big data, including real-time data from the Internet Explore how to use different machine learning models to ask different questions of your data In Detail The Python: Real-World Data Science course will take you on a journey to become an efficient data science practitioner by thoroughly understanding the key concepts of Python. This learning path is divided into four modules and each module are a mini course in their own right, and as you complete each one, you'll have gained key skills and be ready for the material in the next module. The course begins with getting your Python fundamentals nailed down. After getting familiar with Python core concepts, it's time that you dive into the field of data science. In the second module, you'll learn how to perform data analysis using Python in a practical and example-driven way. The third module will teach you how to design and develop data mining applications using a variety of datasets, starting with basic classification and affinity analysis to more complex data types including text, images, and graphs. Machine learning and predictive analytics have become the most important approaches to uncover data gold mines. In the final module, we'll discuss the necessary details regarding machine learning concepts, offering intuitive yet informative explanations on how machine learning algorithms work, how to use them, and most importantly, how to avoid the common pitfalls. Style and approach This course includes all the resources that will help you jump into the data science field with Python and learn how to make sense of data. The aim is to create a smooth learning path that will teach you how to get started with powerful Python libraries and perform various data science techniques in depth.

Summary Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. Without overdosing you on academic theory and complex mathematics, it introduces the day-to-day practice of machine learning, preparing you to successfully build and deploy powerful ML systems. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning systems help you find valuable insights and patterns in data, which you'd never recognize with traditional methods. In the real world, ML techniques give you a way to identify trends, forecast behavior, and make fact-based recommendations. It's a hot and growing field, and up-to-speed ML developers are in demand. About the Book Real-World Machine Learning will teach you the concepts and techniques you need to be a successful machine learning practitioner without overdosing you on abstract theory and complex mathematics. By working through immediately relevant examples in Python, you'll build skills in data acquisition and modeling, classification, and regression. You'll also explore the most important tasks like model validation, optimization, scalability, and real-time streaming. When you're done, you'll be ready to successfully build, deploy, and maintain your own powerful ML systems. What's Inside Predicting future behavior Performance evaluation and optimization Analyzing sentiment and making recommendations About the Reader No prior machine learning experience assumed. Readers should know Python. About the Authors Henrik Brink, Joseph Richards and Mark Fetherolf are experienced data scientists engaged in the daily practice of machine learning. Table of Contents PART 1: THE MACHINE-LEARNING WORKFLOW What is machine learning? Real-world data Modeling and prediction Model evaluation and optimization Basic feature engineering PART 2: PRACTICAL APPLICATION Example: NYC taxi data Advanced feature engineering Advanced NLP example: movie review sentiment Scaling machine-learning workflows Example: digital display advertising

[Copyright: b4ba91630defe96c19daa3c5cec10520](https://www.manning.com/books/python-real-world-data-science)