

## Algorithmic Trading Using Python Dvc Futures

Learn and implement quantitative finance using popular Python libraries like NumPy, pandas, and Keras Key Features Understand Python data structure fundamentals and work with time series data Use popular Python libraries including TensorFlow, Keras, and SciPy to deploy key concepts in quantitative finance Explore various Python programs and learn finance paradigms Book Description Python is one of the most popular languages used for quantitative finance. With this book, you'll explore the key characteristics of Python for finance, solve problems in finance, and understand risk management. The book starts with major concepts and techniques related to quantitative finance, and an introduction to some key Python libraries. Next, you'll implement time series analysis using pandas and DataFrames. The following chapters will help you gain an understanding of how to measure the diversifiable and non-diversifiable security risk of a portfolio and optimize your portfolio by implementing Markowitz Portfolio Optimization. Sections on regression analysis methodology will help you to value assets and understand the relationship between commodity prices and business stocks. In addition to this, you'll be able to forecast stock prices using Monte Carlo simulation. The book will also highlight forecast models that will show you how to determine the price of a call option by analyzing price variation. You'll also use deep learning for financial data analysis and forecasting. In the concluding chapters, you will create neural networks with TensorFlow and Keras for forecasting and prediction. By the end of this book, you will be equipped with the skills you need to perform different financial analysis tasks using Python What you will learn Clean financial data with data preprocessing Visualize financial data using histograms, color plots, and graphs Perform time series analysis with pandas for forecasting Estimate covariance and the correlation between securities and stocks Optimize your portfolio to understand risks when there is a possibility of higher returns Calculate expected returns of a stock to measure the performance of a portfolio manager Create a prediction model using recurrent neural networks (RNN) with Keras and TensorFlow Who this book is for This book is ideal for aspiring data scientists, Python developers and anyone who wants to start performing quantitative finance using Python. You can also make this beginner-level guide your first choice if you're looking to pursue a career as a financial analyst or a data analyst. Working knowledge of Python programming language is necessary.

The definitive deep-dive guide to hardware and software troubleshooting on Cisco Nexus switches The Cisco Nexus platform and NX-OS switch operating system combine to deliver unprecedented speed, capacity, resilience, and flexibility in today's data center networks. Troubleshooting Cisco Nexus Switches and NX-OS is your single reference for quickly identifying and solving problems with these business-critical technologies. Three expert authors draw on deep experience with large Cisco customers, emphasizing the most common issues in real-world deployments, including problems that have caused major data center outages. Their authoritative, hands-on guidance addresses both features and architecture, helping you troubleshoot both control plane forwarding and data plane/data path problems and use NX-OS APIs to automate and simplify troubleshooting. Throughout, you'll find real-world configurations, intuitive illustrations, and practical insights into key platform-specific behaviors. This is an indispensable technical resource for all Cisco network consultants, system/support engineers, network operations professionals, and CCNP/CCIE certification candidates working in the data center domain. · Understand the NX-OS operating system and its powerful troubleshooting tools · Solve problems with cards, hardware drops, fabrics, and CoPP policies · Troubleshoot network packet switching and forwarding · Properly design, implement, and troubleshoot issues related to Virtual Port Channels (VPC and VPC+) · Optimize routing through filtering or path manipulation · Optimize IP/IPv6 services and FHRP protocols (including HSRP, VRRP, and Anycast HSRP) · Troubleshoot EIGRP, OSPF, and IS-IS neighbor relationships and routing paths · Identify and resolve issues with Nexus route maps · Locate problems with BGP neighbor adjacencies and enhance path selection · Troubleshoot high availability components (BFD, SSO, ISSU, and GIR) · Understand multicast protocols and troubleshooting techniques · Identify and solve problems with OTV · Use NX-OS APIs to automate troubleshooting and administrative tasks

Take full creative control of your web applications with Flask, the Python-based microframework. With this hands-on book, you'll learn Flask from the ground up by developing a complete social blogging application step-by-step. Author Miguel Grinberg walks you through the framework's core functionality, and shows you how to extend applications with advanced web techniques such as database migration and web service communication. Rather than impose development guidelines as other frameworks do, Flask leaves the business of extensions up to you. If you have Python experience, this book shows you how to take advantage of that creative freedom. Learn Flask's basic application structure and write an example app Work with must-have components—templates, databases, web forms, and email support Use packages and modules to structure a large application that scales Implement user authentication, roles, and profiles Build a blogging feature by reusing templates, paginating item lists, and working with rich text Use a Flask-based RESTful API to expose app functionality to smartphones, tablets, and other third-party clients Learn how to run unit tests and enhance application performance Explore options for deploying your web app to a production server

Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart programs to access interesting datasets from other web sites, collect data from users of your own applications, and analyze and understand the data once you've found it. Programming Collective Intelligence takes you into the world of machine learning and statistics, and explains how to draw conclusions about user experience, marketing, personal tastes, and human behavior in general -- all from information that you and others collect every day. Each algorithm is described clearly and concisely with code that can immediately be used on your web site, blog, Wiki, or specialized application. This book explains: Collaborative filtering techniques that enable online retailers to recommend products or media Methods of clustering to detect groups of similar items in a large dataset Search engine features -- crawlers, indexers, query engines, and the PageRank algorithm Optimization algorithms that search millions of possible solutions to a problem and choose the best one Bayesian filtering, used in spam filters for classifying documents based on word types and other features Using decision trees not only to make predictions, but to model the way decisions are made Predicting numerical values rather than classifications to build price models Support vector machines to match people in online dating sites Non-negative matrix factorization to find the independent features in a dataset Evolving intelligence for problem solving -- how a computer develops its skill by improving its own code the more it plays a game Each chapter includes exercises for extending the algorithms to make them more powerful. Go beyond simple database-backed applications and put the wealth of Internet data to work for you. "Bravo! I cannot think of a better way for a developer to first learn these algorithms and methods, nor can I think of a better way for me (an old AI dog) to reinvigorate my knowledge of the details." -- Dan Russell, Google "Toby's book does a great job of breaking down the complex subject matter of machine-learning algorithms into practical, easy-to-understand examples that can be directly applied to analysis of social interaction across the Web today. If I had this book two years ago, it would have saved precious time going down some fruitless paths." -- Tim Wolters, CTO, Collective Intellect

Proven techniques for market profile users at any level A "market profile" presents a number of basic elements from the market in an easily understood graphic format that, when analyzed properly, can yield profitable intraday and swing trades that traditional indicators do not reveal. Steidlmayer on Markets shows readers how to find these opportunities using the innovative techniques developed by the author during his many years of trading the market. This fully updated Second Edition covers innovations in both technology and technique-and broadens the scope of "market profile" to include stocks. J. Peter Steidlmayer (Chicago, IL) joined the Chicago Board of Trade in 1963 and has been an independent trader ever since. Steidlmayer served on the Board of Directors of the Board of Trade in 1981-1983. While

a director, he was responsible for initiating Market Profile and the Liquidity Data Bank. Steve Hawkins (Chicago, IL) has experience in trading in both stocks and commodities. Over the past seven years, Hawkins has educated traders across the globe. He has also collaborated on the writing of books on trading and written articles for industry trade publications. He is a graduate of the University of Illinois with a degree in economics. New technology and the advent of around the clock trading have opened the floodgates to both foreign and domestic markets. Traders need the wisdom of industry veterans and the vision of innovators in today's volatile financial marketplace. The Wiley Trading series features books by traders who have survived the market's ever changing temperament and have prospered—some by reinventing systems, others by getting back to basics. Whether a novice trader, professional or somewhere in-between, these books will provide the advice and strategies needed to prosper today and well into the future.

This unified treatment of linear and nonlinear filtering theory presents material previously available only in journals, and in terms accessible to engineering students. Its sole prerequisites are advanced calculus, the theory of ordinary differential equations, and matrix analysis. Although theory is emphasized, the text discusses numerous practical applications as well. Taking the state-space approach to filtering, this text models dynamical systems by finite-dimensional Markov processes, outputs of stochastic difference, and differential equations. Starting with background material on probability theory and stochastic processes, the author introduces and defines the problems of filtering, prediction, and smoothing. He presents the mathematical solutions to nonlinear filtering problems, and he specializes the nonlinear theory to linear problems. The final chapters deal with applications, addressing the development of approximate nonlinear filters, and presenting a critical analysis of their performance.

Algorithmic Trading with Python discusses modern quant trading methods in Python with a heavy focus on pandas, numpy, and scikit-learn. After establishing an understanding of technical indicators and performance metrics, readers will walk through the process of developing a trading simulator, strategy optimizer, and financial machine learning pipeline. This book maintains a high standard of reproducibility. All code and data is self-contained in a GitHub repo. The data includes hyper-realistic simulated price data and alternative data based on real securities. Algorithmic Trading with Python (2020) is the spiritual successor to Automated Trading with R (2016). This book covers more content in less time than its predecessor due to advances in open-source technologies for quantitative analysis.

“We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch

Key Features

- Written by PyTorch’s creator and key contributors
- Develop deep learning models in a familiar Pythonic way
- Use PyTorch to build an image classifier for cancer detection
- Diagnose problems with your neural network and improve training with data augmentation

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About The Book

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It’s great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you’ll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks.

What You Will Learn

- Understanding deep learning data structures such as tensors and neural networks
- Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results
- Implementing modules and loss functions
- Utilizing pretrained models from PyTorch Hub
- Methods for training networks with limited inputs
- Sifting through unreliable results to diagnose and fix problems in your neural network
- Improve your results with augmented data, better model architecture, and fine tuning

This Book Is Written For

For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required.

About The Authors

Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer.

Table of Contents

PART 1 - CORE PYTORCH

- 1 Introducing deep learning and the PyTorch Library
- 2 Pretrained networks
- 3 It starts with a tensor
- 4 Real-world data representation using tensors
- 5 The mechanics of learning
- 6 Using a neural network to fit the data
- 7 Telling birds from airplanes: Learning from images
- 8 Using convolutions to generalize

PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER

- 9 Using PyTorch to fight cancer
- 10 Combining data sources into a unified dataset
- 11 Training a classification model to detect suspected tumors
- 12 Improving training with metrics and augmentation
- 13 Using segmentation to find suspected nodules
- 14 End-to-end nodule analysis, and where to go next

PART 3 - DEPLOYMENT

- 15 Deploying to production

This book includes novel and state-of-the-art research discussions that articulate and report all research aspects, including theoretical and experimental prototypes and applications that incorporate sustainability into emerging applications. In recent years, sustainability and information and communication technologies (ICT) are highly intertwined, where sustainability resources and its management has attracted various researchers, stakeholders, and industrialists. The energy-efficient communication technologies have revolutionized the various smart applications like smart cities, healthcare, entertainment, and business. The book discusses and articulates emerging challenges in significantly reducing the energy consumption of communication systems and also explains development of a sustainable and energy-efficient mobile and wireless communication network. It includes best selected high-quality conference papers in different fields such as internet of things, cloud computing, data mining, artificial intelligence, machine learning, autonomous systems, deep learning, neural networks, renewable energy sources, sustainable wireless communication networks, QoS, network sustainability, and many other related areas.

JavaScript Robotics is on the rise. Rick Waldron, the lead author of this book and creator of the Johnny-Five platform, is at the forefront of this movement. Johnny-Five is an open source JavaScript Arduino programming framework for robotics. This book brings together fifteen innovative programmers, each creating a unique Johnny-Five robot step-by-step, and offering tips and tricks along the way. Experience with JavaScript is a prerequisite.

This book gathers selected papers from Artificial Intelligence and Industrial Applications (A2IA'2020), the first installment of an annual international conference organized by ENSAM-Meknes at Moulay Ismail University, Morocco. The 29 papers presented here were carefully reviewed and selected from 141 submissions by an international scientific committee. They address various aspects of artificial intelligence such as digital twin, multiagent systems, deep learning, image processing and analysis, control, prediction, modeling, optimization and design, as well as AI applications in industry, health, energy, agriculture, and education. The book is intended for AI experts, offering them a valuable overview and global outlook for the future, and highlights a wealth of innovative ideas and recent, important advances in AI applications, both of a foundational and practical nature. It will also appeal to non-experts who are curious about this timely and important subject.

Enhances Python skills by working with data structures and algorithms and gives examples of complex systems using exercises, case studies, and simple explanations.

If you know how to program, you're ready to tackle Bayesian statistics. With this book, you'll learn how to solve statistical problems with Python code instead of mathematical formulas, using discrete

probability distributions rather than continuous mathematics. Once you get the math out of the way, the Bayesian fundamentals will become clearer and you'll begin to apply these techniques to real-world problems. Bayesian statistical methods are becoming more common and more important, but there aren't many resources available to help beginners. Based on undergraduate classes taught by author Allen B. Downey, this book's computational approach helps you get a solid start. Use your programming skills to learn and understand Bayesian statistics Work with problems involving estimation, prediction, decision analysis, evidence, and Bayesian hypothesis testing Get started with simple examples, using coins, dice, and a bowl of cookies Learn computational methods for solving real-world problems This open access book makes a case for a socially inclusive energy transition and illustrates how engineering and public policy professionals can contribute to shaping an inclusive energy transition, building on a socio-technical systems engineering approach. Accomplishing a net-zero greenhouse gas emissions economy in 2050 is a daunting challenge. This book explores the challenges of the energy transition from the perspectives of technological innovation, public policy, social values and ethics. It elaborates on two particular gaps in the design of public policy interventions focused on decarbonization of the energy system and discusses how both could be remedied. First, the siloed organization of public administration fails to account for the many interdependencies between the energy sector, the mobility system, digital infrastructure and the built environment. Cross-sector coordination of policies and policy instruments is needed to avoid potentially adverse effects upon society and the economy, which may hamper the energy transition rather than accelerate it. Second, energy and climate policies pay insufficient attention to the social values at stake in the energy transition. In addressing these gaps, this book intends to inspire decision makers engaged in the energy transition to embrace the transition as an opportunity to bring a more inclusive society into being.

Build Machine Learning models with a sound statistical understanding. About This Book Learn about the statistics behind powerful predictive models with p-value, ANOVA, and F- statistics. Implement statistical computations programmatically for supervised and unsupervised learning through K-means clustering. Master the statistical aspect of Machine Learning with the help of this example-rich guide to R and Python. Who This Book Is For This book is intended for developers with little to no background in statistics, who want to implement Machine Learning in their systems. Some programming knowledge in R or Python will be useful. What You Will Learn Understand the Statistical and Machine Learning fundamentals necessary to build models Understand the major differences and parallels between the statistical way and the Machine Learning way to solve problems Learn how to prepare data and feed models by using the appropriate Machine Learning algorithms from the more-than-adequate R and Python packages Analyze the results and tune the model appropriately to your own predictive goals Understand the concepts of required statistics for Machine Learning Introduce yourself to necessary fundamentals required for building supervised & unsupervised deep learning models Learn reinforcement learning and its application in the field of artificial intelligence domain In Detail Complex statistics in Machine Learning worry a lot of developers. Knowing statistics helps you build strong Machine Learning models that are optimized for a given problem statement. This book will teach you all it takes to perform complex statistical computations required for Machine Learning. You will gain information on statistics behind supervised learning, unsupervised learning, reinforcement learning, and more. Understand the real-world examples that discuss the statistical side of Machine Learning and familiarize yourself with it. You will also design programs for performing tasks such as model, parameter fitting, regression, classification, density collection, and more. By the end of the book, you will have mastered the required statistics for Machine Learning and will be able to apply your new skills to any sort of industry problem. Style and approach This practical, step-by-step guide will give you an understanding of the Statistical and Machine Learning fundamentals you'll need to build models.

In this book readers will find technological discussions on the existing and emerging technologies across the different stages of the big data value chain. They will learn about legal aspects of big data, the social impact, and about education needs and requirements. And they will discover the business perspective and how big data technology can be exploited to deliver value within different sectors of the economy. The book is structured in four parts: Part I "The Big Data Opportunity" explores the value potential of big data with a particular focus on the European context. It also describes the legal, business and social dimensions that need to be addressed, and briefly introduces the European Commission's BIG project. Part II "The Big Data Value Chain" details the complete big data lifecycle from a technical point of view, ranging from data acquisition, analysis, curation and storage, to data usage and exploitation. Next, Part III "Usage and Exploitation of Big Data" illustrates the value creation possibilities of big data applications in various sectors, including industry, healthcare, finance, energy, media and public services. Finally, Part IV "A Roadmap for Big Data Research" identifies and prioritizes the cross-sectorial requirements for big data research, and outlines the most urgent and challenging technological, economic, political and societal issues for big data in Europe. This compendium summarizes more than two years of work performed by a leading group of major European research centers and industries in the context of the BIG project. It brings together research findings, forecasts and estimates related to this challenging technological context that is becoming the major axis of the new digitally transformed business environment.

The go-to guidebook for deploying Big Data solutions withHadoop Today's enterprise architects need to understand how the Hadoopframeworks and APIs fit together, and how they can be integrated todeliver real-world solutions. This book is a practical, detailedguide to building and implementing those solutions, with code-levelinstruction in the popular Wrox tradition. It covers storing datawith HDFS and Hbase, processing data with MapReduce, and automatingdata processing with Oozie. Hadoop security, running Hadoop withAmazon Web Services, best practices, and automating Hadoopprocesses in real time are also covered in depth. With in-depth code examples in Java and XML and the latest onrecent additions to the Hadoop ecosystem, this complete resourcealso covers the use of APIs, exposing their inner workings andallowing architects and developers to better leverage and customizethem. The ultimate guide for developers, designers, and architectswho need to build and deploy Hadoop applications Covers storing and processing data with various technologies,automating data processing, Hadoop security, and deliveringreal-time solutions Includes detailed, real-world examples and code-levelguidelines Explains when, why, and how to use these tools effectively Written by a team of Hadoop experts in theprogrammer-to-programmer Wrox style Professional Hadoop Solutions is the reference enterprisearchitects and developers need to maximize the power of Hadoop.

The volume presents high quality papers presented at the Second International Conference on Microelectronics, Computing & Communication Systems (MCCS 2017). The book discusses recent trends in technology and advancement in MEMS and nanoelectronics, wireless communications, optical communication, instrumentation, signal processing, image processing, bioengineering, green energy, hybrid vehicles, environmental science, weather forecasting, cloud computing, renewable energy, RFID, CMOS sensors, actuators, transducers, telemetry systems, embedded systems, and sensor network applications. It includes original papers based on original theoretical, practical, experimental, simulations, development, application, measurement, and testing. The applications and solutions discussed in the book will serve as a good reference material for future works. The 30-volume set, comprising the LNCS books 12346 until 12375, constitutes the refereed proceedings of the 16th European Conference on Computer Vision, ECCV 2020,

which was planned to be held in Glasgow, UK, during August 23-28, 2020. The conference was held virtually due to the COVID-19 pandemic. The 1360 revised papers presented in these proceedings were carefully reviewed and selected from a total of 5025 submissions. The papers deal with topics such as computer vision; machine learning; deep neural networks; reinforcement learning; object recognition; image classification; image processing; object detection; semantic segmentation; human pose estimation; 3d reconstruction; stereo vision; computational photography; neural networks; image coding; image reconstruction; object recognition; motion estimation.

Become more culturally competent in an increasingly diverse world Recent years have seen dramatic changes to several institutions worldwide. Our increasingly interconnected, digitized, and globalized world presents immense opportunities and unique challenges. Modern businesses and schools interact with individuals and organizations from a diverse range of cultural and national backgrounds—increasing the likelihood for miscommunication, errors in strategy, and unintended consequences in the process. This has also spilled into our daily lives and the way we consume information today. Understanding how to navigate these and other pitfalls requires adaptability, nuanced cross-cultural communication, and effective conflict resolution. *Use Your Difference to Make a Difference* provides readers with a skills-based, actionable plan that transforms differences into agents of inclusiveness, connection, and mutual understanding. This innovative and timely guide illustrates how to leverage differences to move beyond unconscious biases, manage a culturally-diverse workplace, create an environment for more tolerant schooling environments, more trusted media, communicate across borders, find and retain diverse talent, and bridge the gap between working locally and expanding globally. Expert guidance on a comprehensive range of topics—teamwork, leadership styles, information sharing, delegation, supervision, giving and receiving feedback, coaching and motivation, recruiting, managing suppliers and customers, and more—helps you manage the essential aspects of international relationships and cultural awareness. This valuable resource contains the indispensable knowledge required to: Develop self-awareness needed to be a cross-cultural communicator Develop content, messaging techniques, marketing plans, and business strategies that translate across cultural borders Help your employees to better understand and collaborate with clients and colleagues from different backgrounds Help teachers build safe environments for students to be themselves Strengthen cross-cultural competencies in yourself, your team, and your entire organization Understand the cultural, economic, and political factors surrounding our world *Use Your Difference to Make a Difference* is a must-have resource for any educator, parent, leader, manager, or team member of an organization that interacts with co-workers and customers from diverse cultural backgrounds.

The book includes contributions on the latest model-based methods for the development of personal and commercial vehicle control devices. The main topics treated are: application of simulation and model design to development of driver assistance systems; physical and database model design for engines, motors, powertrain, undercarriage and the whole vehicle; new simulation tools, methods and optimization processes; applications of simulation in function and software development; function and software testing using HiL, MiL and SiL simulation; application of simulation and optimization in application of control devices; automation approaches at all stages of the development process.

A new way of thinking about data science and data ethics that is informed by the ideas of intersectional feminism. Today, data science is a form of power. It has been used to expose injustice, improve health outcomes, and topple governments. But it has also been used to discriminate, police, and surveil. This potential for good, on the one hand, and harm, on the other, makes it essential to ask: Data science by whom? Data science for whom? Data science with whose interests in mind? The narratives around big data and data science are overwhelmingly white, male, and techno-heroic. In *Data Feminism*, Catherine D'Ignazio and Lauren Klein present a new way of thinking about data science and data ethics—one that is informed by intersectional feminist thought. Illustrating data feminism in action, D'Ignazio and Klein show how challenges to the male/female binary can help challenge other hierarchical (and empirically wrong) classification systems. They explain how, for example, an understanding of emotion can expand our ideas about effective data visualization, and how the concept of invisible labor can expose the significant human efforts required by our automated systems. And they show why the data never, ever “speak for themselves.” *Data Feminism* offers strategies for data scientists seeking to learn how feminism can help them work toward justice, and for feminists who want to focus their efforts on the growing field of data science. But *Data Feminism* is about much more than gender. It is about power, about who has it and who doesn't, and about how those differentials of power can be challenged and changed.

This book brings together past experience, current work and promising future trends associated with distributed computing, artificial intelligence and their application in order to provide efficient solutions to real problems. DCAI 2020 is a forum to present applications of innovative techniques for studying and solving complex problems in artificial intelligence and computing areas. This years technical program will present both high quality and diversity, with contributions in well-established and evolving areas of research. Specifically, 83 papers were submitted to main track and special sessions, by authors from 26 different countries representing a truly "wide area network" of research activity. The DCAI20 technical program has selected 35 papers and, as in past editions, it will be special issues in ranked journals. This symposium is organized by the University of L'Aquila (Italy). We would like to thank all the contributing authors, the members of the Program Committee and the sponsors (IBM, Armundia Group, EurAI, AEPIA, APPIA, CINI, OIT, UGR, HU, SCU, USAL, AIR Institute and UNIVAQ).

In this book, the author examines the ethical implications of Artificial Intelligence systems as they integrate and replace traditional social structures in new sociocognitive-technological environments. She discusses issues related to the integrity of researchers, technologists, and manufacturers as they design, construct, use, and manage artificially intelligent systems; formalisms for reasoning about moral decisions as part of the behavior of artificial autonomous systems such as agents and robots; and design

methodologies for social agents based on societal, moral, and legal values. Throughout the book the author discusses related work, conscious of both classical, philosophical treatments of ethical issues and the implications in modern, algorithmic systems, and she combines regular references and footnotes with suggestions for further reading. This short overview is suitable for undergraduate students, in both technical and non-technical courses, and for interested and concerned researchers, practitioners, and citizens. Become a master at penetration testing using machine learning with Python Key Features Identify ambiguities and breach intelligent security systems Perform unique cyber attacks to breach robust systems Learn to leverage machine learning algorithms Book Description Cyber security is crucial for both businesses and individuals. As systems are getting smarter, we now see machine learning interrupting computer security. With the adoption of machine learning in upcoming security products, it's important for pentesters and security researchers to understand how these systems work, and to breach them for testing purposes. This book begins with the basics of machine learning and the algorithms used to build robust systems. Once you've gained a fair understanding of how security products leverage machine learning, you'll dive into the core concepts of breaching such systems. Through practical use cases, you'll see how to find loopholes and surpass a self-learning security system. As you make your way through the chapters, you'll focus on topics such as network intrusion detection and AV and IDS evasion. We'll also cover the best practices when identifying ambiguities, and extensive techniques to breach an intelligent system. By the end of this book, you will be well-versed with identifying loopholes in a self-learning security system and will be able to efficiently breach a machine learning system. What you will learn Take an in-depth look at machine learning Get to know natural language processing (NLP) Understand malware feature engineering Build generative adversarial networks using Python libraries Work on threat hunting with machine learning and the ELK stack Explore the best practices for machine learning Who this book is for This book is for pen testers and security professionals who are interested in learning techniques to break an intelligent security system. Basic knowledge of Python is needed, but no prior knowledge of machine learning is necessary.

If you know how to program with Python and also know a little about probability, you're ready to tackle Bayesian statistics. With this book, you'll learn how to solve statistical problems with Python code instead of mathematical notation, and use discrete probability distributions instead of continuous mathematics. Once you get the math out of the way, the Bayesian fundamentals will become clearer, and you'll begin to apply these techniques to real-world problems. Bayesian statistical methods are becoming more common and more important, but not many resources are available to help beginners. Based on undergraduate classes taught by author Allen Downey, this book's computational approach helps you get a solid start. Use your existing programming skills to learn and understand Bayesian statistics Work with problems involving estimation, prediction, decision analysis, evidence, and hypothesis testing Get started with simple examples, using coins, M&Ms, Dungeons & Dragons dice, paintball, and hockey Learn computational methods for solving real-world problems, such as interpreting SAT scores, simulating kidney tumors, and modeling the human microbiome.

This book provides an accessible but concise edited coverage of the main topics, tools and issues in virtual heritage.

Twelve Years a Slave (1853) is a memoir and slave narrative by Solomon Northup, as told to and edited by David Wilson. Northup, a black man who was born free in New York, details his kidnapping in Washington, D.C. and subsequent sale into slavery. After having been kept in bondage for 12 years in Louisiana by various masters, Northup was able to write to friends and family in New York, who were in turn able to secure his release. Northup's account provides extensive details on the slave markets in Washington, D.C. and New Orleans and describes at length cotton and sugar cultivation on major plantations in Louisiana.

This book includes selected papers submitted to the ICADABAI-2017 conference, offering an overview of the new methodologies and presenting innovative applications that are of interest to both academicians and practitioners working in the area of analytics. It discusses predictive analytics applications, machine learning applications, human resource analytics, operations analytics, analytics in finance, methodology and econometric applications. The papers in the predictive analytics applications section discuss web analytics, email marketing, customer churn prediction, retail analytics and sports analytics. The section on machine learning applications then examines healthcare analytics, insurance analytics and machine analytics using different innovative machine learning techniques. Human resource analytics addresses important issues relating to talent acquisition and employability using analytics, while a paper in the section on operations analytics describe an innovative application in oil and gas industry. The papers in the analytics in finance part discuss the use of analytical tools in banking and commodity markets, and lastly the econometric applications part presents interesting banking and insurance applications.

Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

Substantial changes have occurred in the nature of political discourse over the past thirty years. Once, traditional media dominated the political landscape, but in recent years Facebook, Twitter, blogs and Blackberrys have emerged as important tools and platforms for political campaigns. While the Canadian party system has proved surprisingly resilient, the rhythms of political life are now very different. A never-ending 24-hour news cycle has resulted in a never-ending political campaign. The implications of this new political style and its impact on political discourse are issues vigorously debated in this new volume of How Canadians Communicate, as is the question on every politician's mind: How can we draw a generation of digital natives into the current political dialogue? With contributions from such diverse figures as Elly Alboim, Richard Davis, Tom Flanagan, David Marshall, and Roger Epp, How Canadians Communicate IV is the most comprehensive review of political communication in Canada in over three decades - one that poses questions fundamental to the quality of public life. David Taras holds the Ralph

Klein Chair in media studies at Mount Royal University. He served as an expert advisor to the House of Commons Standing Committee on Canadian Heritage and co-edited the first two volumes in the How Canadians Communicate series. He is the co-author of *The Last Word: Media Coverage of the Supreme Court of Canada*. Christopher Waddell is director of the School of Journalism and Communication at Carleton University and holds the Carty Chair in business and financial journalism. He was formerly national editor for *The Globe and Mail* and Parliamentary bureau chief for CBC television news.

This book constitutes the refereed proceedings of the 4th International Conference on Soft Computing, Intelligent Systems, and Information Technology, ICSIIT 2015, held in Bali, Indonesia, in March 2015. The 34 revised full papers presented together with 19 short papers, one keynote and 2 invited talks were carefully reviewed and selected from 92 submissions. The papers cover a wide range of topics related to intelligence in the era of Big Data, such as fuzzy logic and control system; genetic algorithm and heuristic approaches; artificial intelligence and machine learning; similarity-based models; classification and clustering techniques; intelligent data processing; feature extraction; image recognition; visualization techniques; intelligent network; cloud and parallel computing; strategic planning; intelligent applications; and intelligent systems for enterprise, government and society.

Endorsed by all major vendors (Microsoft, Oracle, IBM, and SAP), SOA has quickly become the industry standard for building next-generation software; this practical guide shows readers how to achieve the many benefits of SOA Begins with a look at the architectural principles needed to create successful applications and then goes on to examine the process for designing services and SOA implementations Each stage of the design process has an accompanying chapter that walks readers through the details and provides helpful tips, techniques, and examples The author team of SOA practitioners also provides two unique, comprehensive, end-to-end case studies illustrating the architectural and design techniques presented in the book

Fast Python aggressively rehashes the basics of Python programming in order to dispel myths and misconceptions about how to write fast code. Readers equipped with the lessons from this book will be able to test, diagnose, and optimize out performance bottlenecks in their own work. For each algorithm discussed, readers will walk through numerous progressively faster methods of programming it, all while picking up bits of fundamental knowledge about time complexity, memory efficiency, data structures, multi-threading, and vectorization. As such, this book is relevant to veterans looking refresh their methods and to computer science students navigating Algorithms 101. This book maintains a high standard of reproducibility. All of the graphics, tables, and code profiles contained in this book are fully reproducible and available to anyone in a public GitHub repository.

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learn Cover advanced and state-of-the-art neural network architectures Understand the theory and math behind neural networks Train DNNs and apply them to modern deep learning problems Use CNNs for object detection and image segmentation Implement generative adversarial networks (GANs) and variational autoencoders to generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

Gain a practical introduction to DataOps, a new discipline for delivering data science at scale inspired by practices at companies such as Facebook, Uber, LinkedIn, Twitter, and eBay. Organizations need more than the latest AI algorithms, hottest tools, and best people to turn data into insight-driven action and useful analytical data products. Processes and thinking employed to manage and use data in the 20th century are a bottleneck for working effectively with the variety of data and advanced analytical use cases that organizations have today. This book provides the approach and methods to ensure continuous rapid use of data to create analytical data products and steer decision making. Practical DataOps shows you how to optimize the data supply chain from diverse raw data sources to the final data product, whether the goal is a machine learning model or other data-orientated output. The book provides an approach to eliminate wasted effort and improve collaboration between data producers, data consumers, and the rest of the

organization through the adoption of lean thinking and agile software development principles. This book helps you to improve the speed and accuracy of analytical application development through data management and DevOps practices that securely expand data access, and rapidly increase the number of reproducible data products through automation, testing, and integration. The book also shows how to collect feedback and monitor performance to manage and continuously improve your processes and output.

**What You Will Learn**

- Develop a data strategy for your organization to help it reach its long-term goals
- Recognize and eliminate barriers to delivering data to users at scale
- Work on the right things for the right stakeholders through agile collaboration
- Create trust in data via rigorous testing and effective data management
- Build a culture of learning and continuous improvement through monitoring deployments and measuring outcomes
- Create cross-functional self-organizing teams focused on goals not reporting lines
- Build robust, trustworthy, data pipelines in support of AI, machine learning, and other analytical data products

**Who This Book Is For**

Data science and advanced analytics experts, CIOs, CDOs (chief data officers), chief analytics officers, business analysts, business team leaders, and IT professionals (data engineers, developers, architects, and DBAs) supporting data teams who want to dramatically increase the value their organization derives from data. The book is ideal for data professionals who want to overcome challenges of long delivery time, poor data quality, high maintenance costs, and scaling difficulties in getting data science output and machine learning into customer-facing production.

This book, fully updated for Python version 3.6+, covers the key ideas that link probability, statistics, and machine learning illustrated using Python modules in these areas. All the figures and numerical results are reproducible using the Python codes provided. The author develops key intuitions in machine learning by working meaningful examples using multiple analytical methods and Python codes, thereby connecting theoretical concepts to concrete implementations. Detailed proofs for certain important results are also provided. Modern Python modules like Pandas, Sympy, Scikit-learn, Tensorflow, and Keras are applied to simulate and visualize important machine learning concepts like the bias/variance trade-off, cross-validation, and regularization. Many abstract mathematical ideas, such as convergence in probability theory, are developed and illustrated with numerical examples. This updated edition now includes the Fisher Exact Test and the Mann-Whitney-Wilcoxon Test. A new section on survival analysis has been included as well as substantial development of Generalized Linear Models. The new deep learning section for image processing includes an in-depth discussion of gradient descent methods that underpin all deep learning algorithms. As with the prior edition, there are new and updated \*Programming Tips\* that illustrate effective Python modules and methods for scientific programming and machine learning. There are 445 run-able code blocks with corresponding outputs that have been tested for accuracy. Over 158 graphical visualizations (almost all generated using Python) illustrate the concepts that are developed both in code and in mathematics. We also discuss and use key Python modules such as Numpy, Scikit-learn, Sympy, Scipy, Lifelines, CvxPy, Theano, Matplotlib, Pandas, Tensorflow, Statsmodels, and Keras. This book is suitable for anyone with an undergraduate-level exposure to probability, statistics, or machine learning and with rudimentary knowledge of Python programming.

[Copyright: 949c6b15ed412a75da0040dda461750b](#)