

## Frank Kanes Taming Big Data With Apache Spark And Python

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune, and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

The classic manifesto of the liberated woman, this book explores every facet of a woman's life.

Greek and Roman demigods from the Prophecy of Seven must work together to seal the Doors of Death--and help Percy and Annabeth escape the Underworld in the process.

In American literature, domestic fictions--that is, novels focused on the home and homemaking--are linked with white, middle-class women's fiction and culture. Employing a spatial lens, *Neodomestic American Fiction* joins and extends other studies in redefining domestic fiction's literary history and definition. Unlike previous redefinitions and reevaluations, *Neodomestic American Fiction* reads domestic novels alongside feminist geography and architectural history to map the links and disjunctions among a range of authors writing during the same period as well as across centuries and cultures. Kristin Jacobson's attention to domestic geographies reveals a new space and subgenre emerge in the 1980s: neodomestic fiction. In this innovative study, Kristin Jacobson identifies over thirty novels that renovate traditional forms, therefore challenging model domesticity's conservative gender, racial, and sexual politics. Rather than produce stable single-family homes, neodomestic fictions advance a politics of instability characterized by mobility, renovation and redesign, and relational space. These "alternative" domesticities--when read in the context of neodomestic fiction--are not marginal but rather central to domesticity's configurations. Such resistance, as Iris Marion Young argues, "is integral to modern political theory and is not an alternative to it." Thus, this spatial analysis of post-1980 domestic novels does not indicate a post-feminist or post-gender world. Rather, neodomestic fiction's heterogeneous, unstable spaces offer opportunities to examine contemporary

hierarchies and experiment with more egalitarian homemaking. These fictions include Toni Morrison's *Paradise*, Barbara Kingsolver's *The Poisonwood Bible*, Leslie Marmon Silko's *Gardens in the Dunes*, and Chang-rae Lee's *A Gesture Life*.

Contains highlights of an American College of Neuropsychopharmacology conference on new directions in the development of atypical and other novel antipsychotic drugs. Presents new theories and preclinical and clinical data on various drugs and classes of drugs including amperozide and other drugs.

Ku and Hina-man and woman-were the great ancestral gods of heaven and earth for the ancient Hawaiians. They were life's fruitfulness and all the generations of mankind, both those who are to come and those already born. The Hawaiian gods were like great chiefs from far lands who visited among the people, entering their daily lives sometimes as humans or animals, sometimes taking residence in a stone or wooden idol. As years passed, the families of gods grew and included the trickster Maui, who snared the sun, and fiery Pele of the volcano. Ancient Hawaiian lived by the animistic philosophy that assigned living souls to animals, trees, stones, stars, and clouds, as well as to humans. Religion and mythology were interwoven in Hawaiian culture; and local legends and genealogies were preserved in song, chant, and narrative. Martha Beckwith was the first scholar to chart a path through the hundreds of books, articles, and little-known manuscripts that recorded the oral narratives of the Hawaiian people. Her book has become a classic work of folklore and ethnology, and the definitive treatment of Hawaiian mythology. With an introduction by Katherine Luomala.

*French Feminism Reader* is a collection of essays representing the authors and issues from French theory most influential in the American context. The book is designed for use in courses, and it includes illuminating introductions to the work of each author. These introductions include biographical information, influences and intellectual context, major themes in the author's work as a whole, and specific introductions to the selections in this volume. The contributors represent the two trends in French theory that have proven most useful to American feminists: social theory and psychoanalytic theory. Both of these trends move away from any traditional discussions of nature toward discussions of socially constructed notions of sex, sexuality and gender roles. While feminists interested in social theory focus on the ways in which social institutions shape these notions, feminists interested in psychoanalytic theory focus on cultural representations of sex, sexuality and gender roles, and the ways that they affect the psyche. This collection includes selections by Simone de Beauvoir, Christine Delphy, Colette Guillaumin, Monique Wittig, Michele Le Doeuff, Julia Kristeva, Luce Irigaray, and Helene Cixous.

Unlike some other reproductions of classic texts (1) We have not used OCR (Optical Character Recognition), as this leads to bad quality books with introduced typos. (2) In books where there are images such as portraits, maps, sketches etc We have endeavoured to keep the quality of these images, so they

represent accurately the original artefact. Although occasionally there may be certain imperfections with these old texts, we feel they deserve to be made available for future generations to enjoy.

Frank Kane's hands-on Spark training course, based on his bestselling Taming Big Data with Apache Spark and Python video, now available in a book. Understand and analyze large data sets using Spark on a single system or on a cluster. About This Book Understand how Spark can be distributed across computing clusters Develop and run Spark jobs efficiently using Python A hands-on tutorial by Frank Kane with over 15 real-world examples teaching you Big Data processing with Spark Who This Book Is For If you are a data scientist or data analyst who wants to learn Big Data processing using Apache Spark and Python, this book is for you. If you have some programming experience in Python, and want to learn how to process large amounts of data using Apache Spark, Frank Kane's Taming Big Data with Apache Spark and Python will also help you. What You Will Learn Find out how you can identify Big Data problems as Spark problems Install and run Apache Spark on your computer or on a cluster Analyze large data sets across many CPUs using Spark's Resilient Distributed Datasets Implement machine learning on Spark using the MLlib library Process continuous streams of data in real time using the Spark streaming module Perform complex network analysis using Spark's GraphX library Use Amazon's Elastic MapReduce service to run your Spark jobs on a cluster In Detail Frank Kane's Taming Big Data with Apache Spark and Python is your companion to learning Apache Spark in a hands-on manner. Frank will start you off by teaching you how to set up Spark on a single system or on a cluster, and you'll soon move on to analyzing large data sets using Spark RDD, and developing and running effective Spark jobs quickly using Python. Apache Spark has emerged as the next big thing in the Big Data domain – quickly rising from an ascending technology to an established superstar in just a matter of years. Spark allows you to quickly extract actionable insights from large amounts of data, on a real-time basis, making it an essential tool in many modern businesses. Frank has packed this book with over 15 interactive, fun-filled examples relevant to the real world, and he will empower you to understand the Spark ecosystem and implement production-grade real-time Spark projects with ease. Style and approach Frank Kane's Taming Big Data with Apache Spark and Python is a hands-on tutorial with over 15 real-world examples carefully explained by Frank in a step-by-step manner. The examples vary in complexity, and you can move through them at your own pace.

Discover the capabilities of PySpark and its application in the realm of data science. This comprehensive guide with hand-picked examples of daily use cases will walk you through the end-to-end predictive model-building cycle with the latest techniques and tricks of the trade. Applied Data Science Using PySpark is divided into six sections which walk you through the book. In section 1, you start with the basics of PySpark focusing on data manipulation. We make you comfortable with the language and then build upon it to introduce you to the mathematical functions available off the shelf. In section 2, you will dive into the art of variable selection where we demonstrate various selection techniques available in PySpark. In section 3, we take you on a journey through machine learning algorithms, implementations, and fine-tuning techniques. We will also talk about different validation metrics and how to use them for picking the best models. Sections 4 and 5 go through machine learning pipelines and various methods available to operationalize the model and serve it through Docker/an API. In the final section, you will cover reusable objects for easy experimentation and learn some tricks that can help you optimize your programs and machine learning pipelines. By the end of this book, you will have seen the flexibility and advantages of PySpark in data science applications. This book is recommended to those who want to unleash the power of parallel computing by simultaneously working with big datasets. What You Will Learn Build an end-to-end predictive

model Implement multiple variable selection techniques Operationalize models Master multiple algorithms and implementations Who This Book is For Data scientists and machine learning and deep learning engineers who want to learn and use PySpark for real-time analysis of streaming data.

The travel experience filled with personal trauma; the pilgrimage through a war-torn place; the journey with those suffering: these represent the darker sides of travel. This book explores the writings and texts of dark journeys and travels. In traveling over the dead, amongst the dying, and alongside the suffering, the authors give us a tour of humanity's violence and misery. From this dark side, there comes great beauty and poignancy in the characterization of flight; creativity in the comic, graphic, and graffiti sketches and comments on life; and the sense of profound and spiritual journeys being undertaken, recorded, and memorialized.

A completely updated, revised edition of the classic, outfitted with a whole new arsenal of indispensable knowledge on global affairs, popular culture, economic trends, scientific principles, and modern arts. Here's your chance to brush up on all those subjects you slept through in school, reacquaint yourself with all the facts you once knew (then promptly forgot), catch up on major developments in the world today, and become the Renaissance man or woman you always knew you could be! How do you tell the Balkans from the Caucasus? What's the difference between fission and fusion? Whigs and Tories? Shiites and Sunnis? Deduction and induction? Why aren't all Shakespearean comedies necessarily thigh-slappers? What are transcendental numbers and what are they good for? What really happened in Plato's cave? Is postmodernism dead or just having a bad hair day? And for extra credit, when should you use the adjective continual and when should you use continuous? An Incomplete Education answers these and thousands of other questions with incomparable wit, style, and clarity. American Studies, Art History, Economics, Film, Literature, Music, Philosophy, Political Science, Psychology, Religion, Science, and World History: Here's the bottom line on each of these major disciplines, distilled to its essence and served up with consummate flair. In this revised edition you'll find a vitally expanded treatment of international issues, reflecting the seismic geopolitical upheavals of the past decade, from economic free-fall in South America to Central Africa's world war, and from violent radicalization in the Muslim world to the crucial trade agreements that are defining globalization for the twenty-first century. And don't forget to read the section "A Nervous American's Guide to Living and Loving on Five Continents" before you answer a personal ad in the International Herald Tribune. As delightful as it is illuminating, An Incomplete Education packs ten thousand years of culture into a single superbly readable volume. This is a book to celebrate, to share, to give and receive, to pore over and browse through, and to return to again and again.

"Like man, woman is a human being." When *The Second Sex* was first published in Paris in 1949--groundbreaking, risqué, brilliantly written and strikingly modern--it provoked both outrage and inspiration. *The Independent Woman* contains three key chapters of Beauvoir's masterwork, which illuminate the feminine condition and identify practical social reforms for gender equality. It captures the essence of the spirited manifesto that switched on light bulbs in the heads of a generation of women and continues to exert profound influence on feminists today.

A concise guide to implementing Spark Big Data analytics for Python developers, and building a real-time and insightful trend tracker data intensive app About This Book • Set up real-time streaming and batch data intensive infrastructure using Spark and Python • Deliver insightful visualizations in a web app using Spark (PySpark) • Inject live data using Spark Streaming with real-time events Who This Book Is For This book is for data scientists and software developers with a focus on Python who want to work with the Spark engine, and it will also benefit Enterprise Architects. All you need to have is a good background of Python and an inclination to work with Spark. What You Will Learn • Create a Python development environment powered

by Spark (PySpark), Blaze, and Bokeh• Build a real-time trend tracker data intensive app• Visualize the trends and insights gained from data using Bokeh• Generate insights from data using machine learning through Spark MLlib• Juggle with data using Blaze• Create training data sets and train the Machine Learning models• Test the machine learning models on test datasets• Deploy the machine learning algorithms and models and scale it for real-time events

In Detail Looking for a cluster computing system that provides high-level APIs? Apache Spark is your answer—an open source, fast, and general purpose cluster computing system. Spark's multi-stage memory primitives provide performance up to 100 times faster than Hadoop, and it is also well-suited for machine learning algorithms. Are you a Python developer inclined to work with Spark engine? If so, this book will be your companion as you create data-intensive app using Spark as a processing engine, Python visualization libraries, and web frameworks such as Flask. To begin with, you will learn the most effective way to install the Python development environment powered by Spark, Blaze, and Bokeh. You will then find out how to connect with data stores such as MySQL, MongoDB, Cassandra, and Hadoop. You'll expand your skills throughout, getting familiarized with the various data sources (Github, Twitter, Meetup, and Blogs), their data structures, and solutions to effectively tackle complexities. You'll explore datasets using iPython Notebook and will discover how to optimize the data models and pipeline. Finally, you'll get to know how to create training datasets and train the machine learning models. By the end of the book, you will have created a real-time and insightful trend tracker data-intensive app with Spark.

**Style and approach** This is a comprehensive guide packed with easy-to-follow examples that will take your skills to the next level and will get you up and running with Spark.

"Saints, Scholars, and Schizophrenics, in its original form--now integrally reproduced in the new edition--is a most important seminal study of an Irish community."—Conor Cruise O'Brien

This book examines Hittite religion from a historical point of view, stressing two basically different stages in its development. The Old Hittite pantheon of the capital Hattusa maintains the indigenous religious tradition of the Hattians without any trace of Mesopotamian, Hurrian or Syrian influence, although Hittite and Luwian deities were worshiped in the family and house cults. The Hittite religion of the Empire period has been examined from a new viewpoint. At the time there were two official pantheons in the state and the dynastic cult respectively. The former is an amalgam of Hittite, Luwian, Hurrian, Syrian and Mesopotamian deities organized on a geographical principle, whereas the latter is purely Hurrian, reflecting the religious beliefs of the new royal family of Kizzuwatnan origin that also influenced local pantheons of central and northern Anatolia. Through the Hurrians, Mesopotamian and Syrian cults were adopted. Simultaneously, many aspects of the Luwian religious tradition were absorbed into both the state and local cults.

Despite the recent turn to affects and emotions in the humanities and despite the unceasing popularity of romantic and erotic love as a motif in fictional works of all genres, the subject has received surprisingly little attention in academic studies of contemporary drama. *Love in Contemporary British Drama* reflects the appeal of love as a topic and driving force in dramatic works with in-depth analyses of eight pivotal plays from the past three decades. Following an interdisciplinary and historical approach, the study collects and condenses theories of love from philosophy and sociology to derive persisting discourses and to examine their reoccurrence and transformation in contemporary plays. Special emphasis is put on narratives of love's compensatory function and precariousness and on how modifications of these narratives epitomise the peculiarities of emotional life in the social and cultural context of the present. Based on the assumption that

drama is especially inclined to draw on shared narratives for representations of love, the book demonstrates that love is both a window to remnants of the past in the present and a proper subject matter for drama in times in which the suitability of the dramatic form has been questioned.

A comprehensive guide to the scope of contemporary urban design theory in Europe and the USA.

Contributions by Timothy P. Barnard, Michael Cohen, Rayna Denison, Martin Flanagan, Sophie Geoffroy-Menoux, Mel Gibson, Kerry Gough, Jonathan Gray, Craig Hight, Derek Johnson, Pascal Lefevre, Paul M. Malone, Neil Rae, Aldo J. Regalado, Jan van der Putten, and David Wilt In Film and Comic Books contributors analyze the problems of adapting one medium to another; the translation of comics aesthetics into film; audience expectations, reception, and reaction to comic book-based films; and the adaptation of films into comics. A wide range of comic/film adaptations are explored, including superheroes (Spider-Man), comic strips (Dick Tracy), realist and autobiographical comics (American Splendor; Ghost World), and photo-montage comics (Mexico's El Santo). Essayists discuss films beginning with the 1978 Superman. That success led filmmakers to adapt a multitude of comic books for the screen including Marvel's Uncanny X-Men, the Amazing Spider-Man, Blade, and the Incredible Hulk as well as alternative graphic novels such as From Hell, V for Vendetta, and Road to Perdition. Essayists also discuss recent works from Mexico, France, Germany, and Malaysia.

Now revised and updated to incorporate numerous new materials, this is the major source for researching American Christian activity in China, especially that of missions and missionaries. It provides a thorough introduction and guide to primary and secondary sources on Christian enterprises and individuals in China that are preserved in hundreds of libraries, archives, historical societies, headquarters of religious orders, and other repositories in the United States. It includes data from the beginnings of Christianity in China in the early eighth century through 1952, when American missionary activity in China virtually ceased. For this new edition, the institutional base has shifted from the Princeton Theological Seminary (Protestant) to the Ricci Institute for Chinese-Western Cultural Relations at the University of San Francisco (Jesuit), reflecting the ecumenical nature of this monumental undertaking.

Jean Baudrillard's classic text was one of the first to focus on the process and meaning of consumption in contemporary culture. Originally published in 1970, the book makes a vital contribution to current debates on consumption. The book includes Baudrillard's most organized discussion of mass media culture, the meaning of leisure, and anomie in affluent society. A chapter on the body demonstrates Baudrillard's extraordinary prescience for flagging vital subjects in contemporary culture long before others. This English translation begins with a new introductory essay.

Consisting of sixteen original essays by experts in the field, including leading and

lesser-known international scholars, *Global Frankenstein* considers the tremendous adaptability and rich afterlives of Mary Shelley's iconic novel, *Frankenstein*, at its bicentenary, in such fields and disciplines as digital technology, film, theatre, dance, medicine, book illustration, science fiction, comic books, science, and performance art. This ground-breaking, celebratory volume, edited by two established Gothic Studies scholars, reassesses *Frankenstein's* global impact for the twenty-first century across a myriad of cultures and nations, from Japan, Mexico, and Turkey, to Britain, Iraq, Europe, and North America. Offering compelling critical dissections of reincarnations of *Frankenstein*, a generically hybrid novel described by its early reviewers as a "bold," "bizarre," and "impious" production by a writer "with no common powers of mind", this collection interrogates its sustained relevance over two centuries during which it has engaged with such issues as mortality, global capitalism, gender, race, embodiment, neoliberalism, disability, technology, and the role of science.

In this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—classification, collaborative filtering, and anomaly detection among others—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find these patterns useful for working on your own data applications. Patterns include:

- Recommending music and the Audioscrobbler data set
- Predicting forest cover with decision trees
- Anomaly detection in network traffic with K-means clustering
- Understanding Wikipedia with Latent Semantic Analysis
- Analyzing co-occurrence networks with GraphX
- Geospatial and temporal data analysis on the New York City Taxi Trips data
- Estimating financial risk through Monte Carlo simulation
- Analyzing genomics data and the BDG project
- Analyzing neuroimaging data with PySpark and Thunder

Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark, and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark. Learn about DataFrames, SQL, and Datasets—Spark's core APIs—through worked examples. Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames. Understand how Spark

runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation

No need to spend hours ploughing through endless data – let Spark, one of the fastest big data processing engines available, do the hard work for you. Key Features Get up and running with Apache Spark and Python Integrate Spark with AWS for real-time analytics Apply processed data streams to machine learning APIs of Apache Spark Book Description Processing big data in real time is challenging due to scalability, information consistency, and fault-tolerance. This book teaches you how to use Spark to make your overall analytical workflow faster and more efficient. You'll explore all core concepts and tools within the Spark ecosystem, such as Spark Streaming, the Spark Streaming API, machine learning extension, and structured streaming. You'll begin by learning data processing fundamentals using Resilient Distributed Datasets (RDDs), SQL, Datasets, and Dataframes APIs. After grasping these fundamentals, you'll move on to using Spark Streaming APIs to consume data in real time from TCP sockets, and integrate Amazon Web Services (AWS) for stream consumption. By the end of this book, you'll not only have understood how to use machine learning extensions and structured streams but you'll also be able to apply Spark in your own upcoming big data projects. What you will learn Write your own Python programs that can interact with Spark Implement data stream consumption using Apache Spark Recognize common operations in Spark to process known data streams Integrate Spark streaming with Amazon Web Services (AWS) Create a collaborative filtering model with the movielens dataset Apply processed data streams to Spark machine learning APIs Who this book is for Data Processing with Apache Spark is for you if you are a software engineer, architect, or IT professional who wants to explore distributed systems and big data analytics. Although you don't need any knowledge of Spark, prior experience of working with Python is recommended.

Learn to build web applications with Catalyst, the popular open source web framework based on the Perl programming language. The Definitive Guide to Catalyst: Writing Extendable, Scalable, and Maintainable Perl-Based Web Applications is a definitive guide to Catalyst version 5.8. This book contains Training materials for new and experience programmers. Worked examples and cookbook-style recipes of common web application programming tasks Fundamentals of web application design and best-practice application style Build data-intensive applications locally and deploy at scale using the combined powers of Python and Spark 2.0 About This Book Learn why and how you can efficiently use Python to process data and build machine learning models in Apache Spark 2.0 Develop and deploy efficient, scalable real-time Spark solutions Take your understanding of using Spark with Python to the next level with this jump start guide Who This Book Is For If you are a Python developer who wants to learn about the Apache Spark 2.0 ecosystem, this book is for you. A firm understanding of Python is expected to get the best out of the book. Familiarity with Spark would be useful, but is not mandatory. What You Will Learn Learn about Apache Spark and the Spark 2.0 architecture Build and interact with Spark DataFrames using Spark SQL Learn how to solve graph and deep learning problems using GraphFrames and TensorFrames respectively Read, transform, and understand data and use it to train machine learning models Build machine learning models with MLlib and ML Learn how to submit your applications programmatically using spark-submit Deploy locally built applications to a cluster In Detail Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. This book will show you how to leverage the power of Python and put it to use in the Spark ecosystem. You will start by getting a firm understanding of the Spark 2.0 architecture and how to set up a Python environment for Spark.



You will get familiar with the modules available in PySpark. You will learn how to abstract data with RDDs and DataFrames and understand the streaming capabilities of PySpark. Also, you will get a thorough overview of machine learning capabilities of PySpark using ML and MLlib, graph processing using GraphFrames, and polyglot persistence using Blaze. Finally, you will learn how to deploy your applications to the cloud using the spark-submit command. By the end of this book, you will have established a firm understanding of the Spark Python API and how it can be used to build data-intensive applications. Style and approach This book takes a very comprehensive, step-by-step approach so you understand how the Spark ecosystem can be used with Python to develop efficient, scalable solutions. Every chapter is standalone and written in a very easy-to-understand manner, with a focus on both the hows and the whys of each concept.

The first volume of its kind to bring together generous selections of the works of three of the great Stoic philosophers, Seneca, Epictetus, and Marcus Aurelius.

Neela Kane worked as a high school teacher for 28 years. She taught children of all ages from toddlers to teenagers. She has written and directed plays for children and won prizes for best dramatics in school competitions. She has a Master's degree in Education.

Quickly find solutions to common programming problems encountered while processing big data. Content is presented in the popular problem-solution format. Look up the programming problem that you want to solve. Read the solution. Apply the solution directly in your own code. Problem solved! PySpark Recipes covers Hadoop and its shortcomings. The architecture of Spark, PySpark, and RDD are presented. You will learn to apply RDD to solve day-to-day big data problems. Python and NumPy are included and make it easy for new learners of PySpark to understand and adopt the model. What You Will Learn Understand the advanced features of PySpark2 and SparkSQL Optimize your code Program SparkSQL with Python Use Spark Streaming and Spark MLlib with Python Perform graph analysis with GraphFrames Who This Book Is For Data analysts, Python programmers, big data enthusiasts

Learn how to build recommender systems from one of Amazon's pioneers in the field. Frank Kane spent over nine years at Amazon, where he managed and led the development of many of Amazon's personalized product recommendation technologies. You've seen automated recommendations everywhere - on Netflix's home page, on YouTube, and on Amazon as these machine learning algorithms learn about your unique interests, and show the best products or content for you as an individual. These technologies have become central to the largest, most prestigious tech employers out there, and by understanding how they work, you'll become very valuable to them. This book is adapted from Frank's popular online course published by Sundog Education, so you can expect lots of visual aids from its slides and a conversational, accessible tone throughout the book. The graphics and scripts from over 300 slides are included, and you'll have access to all of the source code associated with it as well. We'll cover tried and true recommendation algorithms based on neighborhood-based collaborative filtering, and work our way up to more modern techniques including matrix factorization and even deep learning with artificial neural networks. Along the way, you'll learn from Frank's extensive industry experience to understand the real-world challenges you'll encounter when applying these algorithms at large scale and with real-world data. This book is very hands-on; you'll develop your own framework for evaluating and combining many different recommendation algorithms together, and you'll even build your own neural networks using Tensorflow to generate recommendations from real-world movie ratings from real people. We'll cover:

- Building a recommendation engine
- Evaluating recommender systems
- Content-based filtering using item attributes
- Neighborhood-based collaborative filtering with user-based, item-based, and KNN CF
- Model-based methods including matrix factorization and SVD
- Applying deep learning, AI, and artificial neural networks to recommendations
- Session-based recommendations with recursive neural networks
- Scaling to massive data sets with Apache

Spark machine learning, Amazon DSSTNE deep learning, and AWS SageMaker with factorization machines-Real-world challenges and solutions with recommender systems-Case studies from YouTube and Netflix-Building hybrid, ensemble recommenders This comprehensive book takes you all the way from the early days of collaborative filtering, to bleeding-edge applications of deep neural networks and modern machine learning techniques for recommending the best items to every individual user. The coding exercises for this book use the Python programming language. We include an intro to Python if you're new to it, but you'll need some prior programming experience in order to use this book successfully. We also include a short introduction to deep learning, Tensorflow, and Keras if you are new to the field of artificial intelligence, but you'll need to be able to understand new computer algorithms. Dive in, and learn about one of the most interesting and lucrative applications of machine learning and deep learning there is!

In the tradition of Dorothea Lange and Robert Frank, an eye-opening portrait of the rise and fall of the American working class, and a shockingly intimate visual history of Troy, New York that arcs over five hundred years—from Henry Hudson to the industrial revolution to a group of contemporary young women as they grow, survive, and love. Welcome to Troy, New York. The land where mastodon roamed, the Mohicans lived, and the Dutch settled in the seventeenth century. Troy grew from a small trading post into a jewel of the Industrial Revolution. Horseshoes, rail ties, and detachable shirt collars were made there and the middle class boomed, making Troy the fourth wealthiest city per capita in the country. Then, the factories closed, the middle class disappeared, and the downtown fell into disrepair. Troy is the home of Uncle Sam, the Rensselaer Polytechnic Institute, the Rensselaer County Jail, the photographer Brenda Ann Kenneally, and the small group of young women, their children, lovers, and families who Kenneally has been photographing for over a decade. Before Kenneally left Troy, her life looked a lot like the lives of these girls. With passion and profound empathy she has chronicled three generations—their love and heartbreak; their births and deaths; their struggles with poverty, with education, and with each other; and their joy. Brenda Ann Kenneally is the Dorothea Lange of our time—her work a bridge between the people she photographs, history, and us. What began as a brief assignment for The New York Times Magazine became an eye-opening portrait of the rise and fall of the American working class, and a shockingly intimate visual history of Troy that arcs over five hundred years. Kenneally beautifully layers archival images with her own photographs and collages to depict the transformations of this quintessentially American city. The result is a profound, powerful, and intimate look at America, at poverty, at the shrinking middle class, and of people as they grow, survive, and love.

Combine the power of Apache Spark and Python to build effective big data applications

### Key Features

Perform effective data processing, machine learning, and analytics using PySpark  
Overcome challenges in developing and deploying Spark solutions using Python  
Explore recipes for efficiently combining Python and Apache Spark to process data

### Book Description

Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. The PySpark Cookbook presents effective and time-saving recipes for leveraging the power of Python and putting it to use in the Spark ecosystem. You'll start by learning the Apache Spark architecture and how to set up a Python environment for Spark. You'll then get familiar with the modules available in PySpark and start using them effortlessly. In addition to this, you'll discover how to abstract data with RDDs and DataFrames, and understand the streaming capabilities of PySpark. You'll then move on to using ML and MLlib in order to solve any problems related to the machine learning capabilities of PySpark and use GraphFrames to solve graph-processing problems. Finally, you will explore how to deploy your applications to the cloud using the spark-submit command.

By the end of this book, you will be able to use the Python API for Apache Spark to solve any problems associated with building data-intensive applications. What you will learn Configure a local instance of PySpark in a virtual environment Install and configure Jupyter in local and multi-node environments Create DataFrames from JSON and a dictionary using pyspark.sql Explore regression and clustering models available in the ML module Use DataFrames to transform data used for modeling Connect to PubNub and perform aggregations on streams Who this book is for The PySpark Cookbook is for you if you are a Python developer looking for hands-on recipes for using the Apache Spark 2.x ecosystem in the best possible way. A thorough understanding of Python (and some familiarity with Spark) will help you get the best out of the book.

The Advances in Pharmacology series presents a variety of chapters from the best authors in the field. Includes the authority and expertise of leading contributors in pharmacology Presents the latest release in the Advances in Pharmacology series The XXX Filmography, 1968-1988 features more than 3,000 A to Z entries, covering the historic, artistic, and technical aspects of adult cinema from those years, including 35mm features, 16mm storefront features, and 8mm loops. It provides director, producer, cast, screenwriter, cinematographer, and composer listings, with a detailed synopsis for each film. Production company credits, release dates and running times are also given. A DVD appendix lists all titles currently available on DVD, and complete cast and director indexes make this work the most comprehensive guide to Golden Age triple-X films ever published.

This book covers the fundamentals of machine learning with Python in a concise and dynamic manner. It covers data mining and large-scale machine learning using Apache Spark. About This Book Take your first steps in the world of data science by understanding the tools and techniques of data analysis Train efficient Machine Learning models in Python using the supervised and unsupervised learning methods Learn how to use Apache Spark for processing Big Data efficiently Who This Book Is For If you are a budding data scientist or a data analyst who wants to analyze and gain actionable insights from data using Python, this book is for you. Programmers with some experience in Python who want to enter the lucrative world of Data Science will also find this book to be very useful, but you don't need to be an expert Python coder or mathematician to get the most from this book. What You Will Learn Learn how to clean your data and ready it for analysis Implement the popular clustering and regression methods in Python Train efficient machine learning models using decision trees and random forests Visualize the results of your analysis using Python's Matplotlib library Use Apache Spark's MLlib package to perform machine learning on large datasets In Detail Join Frank Kane, who worked on Amazon and IMDb's machine learning algorithms, as he guides you on your first steps into the world of data science. Hands-On Data Science and Python Machine Learning gives you the tools that you need to understand and explore the core topics in the field, and the confidence and practice to build and analyze your own machine learning models. With the help of interesting and easy-to-follow practical examples, Frank Kane explains potentially complex topics such as Bayesian methods and K-means clustering in a way that anybody can understand them. Based on Frank's successful data science course, Hands-On Data Science and Python Machine Learning empowers you to conduct data analysis and perform efficient

machine learning using Python. Let Frank help you unearth the value in your data using the various data mining and data analysis techniques available in Python, and to develop efficient predictive models to predict future results. You will also learn how to perform large-scale machine learning on Big Data using Apache Spark. The book covers preparing your data for analysis, training machine learning models, and visualizing the final data analysis. Style and approach This comprehensive book is a perfect blend of theory and hands-on code examples in Python which can be used for your reference at any time.

[Copyright: ca43f098e476ba0367876876ca954ab1](#)