

---

# Solr In Action Meap Chapter 1

---

A guide for managers and the rest of us  
 Beyond static sites with JavaScript, APIs, and markup  
 How to Find, Organize, and Manipulate It  
 React in Action  
 Microservices Patterns  
 Classical Mechanics  
 Play Framework Cookbook  
 Apache Pulsar in Action  
 A Brain-Friendly Guide  
 Mahout in Action  
 Spring Boot in Action  
 CMIS and Apache Chemistry in Action  
 Covers MongoDB version 3.0  
 Kubernetes in Action  
 Relevant Search  
 Spring Roo in Action  
 Enterprise OSGi In Action  
 A Comprehensive Guide to Building Real-World NLP Systems  
 Grokking Machine Learning  
 Covers Apache Spark 3 with Examples in Java, Python, and Scala  
 Hadoop in Action  
 Real-World Natural Language Processing  
 HBase in Action  
 With Suspense and Concurrent Mode  
 PostGIS in Action, Third Edition  
 Solr in Action  
 Human-in-the-Loop Machine Learning  
 Making Sense of NoSQL  
 Lucene in Action  
 The Jamstack Book  
 Deep Learning for Search  
 Spark in Action  
 With examples in Java  
 Practical Natural Language Processing  
 Practical applications with deep learning  
 Ant in Action  
 PHP in Action  
 Deep Learning with Python, Second Edition  
 React Hooks in Action  
 Graph-Powered Machine Learning

*Solr In Action Meap Chapter 1*

Downloaded from [ftp.wtvq.com](http://ftp.wtvq.com) by guest

---

## PAGE OCONNELL

---

**A guide for managers and the rest of us** Simon and Schuster  
 Graph Databases in Action teaches readers everything they need to know to begin building and running applications powered by graph databases. Right off the bat, seasoned graph database experts introduce readers to just enough graph theory, the graph database ecosystem, and a variety of datastores. They also explore modelling basics in action with real-world examples, then go hands-on with querying, coding traversals, parsing results, and other essential tasks as readers build their own graph-backed social network app complete with a recommendation engine! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.  
Beyond static sites with JavaScript, APIs, and markup Simon and Schuster  
 Hadoop in Action teaches readers how to use Hadoop and write MapReduce programs. The intended readers are programmers, architects, and project managers who have to process large amounts of data offline. Hadoop in Action will lead the reader

from obtaining a copy of Hadoop to setting it up in a cluster and writing data analytic programs. The book begins by making the basic idea of Hadoop and MapReduce easier to grasp by applying the default Hadoop installation to a few easy-to-follow tasks, such as analyzing changes in word frequency across a body of documents. The book continues through the basic concepts of MapReduce applications developed using Hadoop, including a close look at framework components, use of Hadoop for a variety of data analysis tasks, and numerous examples of Hadoop in action. Hadoop in Action will explain how to use Hadoop and present design patterns and practices of programming MapReduce. MapReduce is a complex idea both conceptually and in its implementation, and Hadoop users are challenged to learn all the knobs and levers for running Hadoop. This book takes you beyond the mechanics of running Hadoop, teaching you to write meaningful programs in a MapReduce framework. This book assumes the reader will have a basic familiarity with Java, as most code examples will be written in Java. Familiarity with basic statistical concepts (e.g. histogram, correlation) will help the reader appreciate the more advanced data processing examples. Purchase of the print book comes with an offer of a free PDF,

ePub, and Kindle eBook from Manning. Also available is all code from the book.

How to Find, Organize, and Manipulate It Simon and Schuster

Summary Mahout in Action is a hands-on introduction to machine learning with Apache Mahout. Following real-world examples, the book presents practical use cases and then illustrates how Mahout can be applied to solve them. Includes a free audio- and video-enhanced ebook. About the Technology A computer system that learns and adapts as it collects data can be really powerful. Mahout, Apache's open source machine learning project, captures the core algorithms of recommendation systems, classification, and clustering in ready-to-use, scalable libraries. With Mahout, you can immediately apply to your own projects the machine learning techniques that drive Amazon, Netflix, and others. About this Book This book covers machine learning using Apache Mahout. Based on experience with real-world applications, it introduces practical use cases and illustrates how Mahout can be applied to solve them. It places particular focus on issues of scalability and how to apply these techniques against large data sets using the Apache Hadoop framework. This book is written for developers familiar with Java -- no prior experience with Mahout is assumed. Owners of a Manning pBook purchased anywhere in the world can download a free eBook from manning.com at any time. They can do so multiple times and in any or all formats available (PDF, ePub or Kindle). To do so, customers must register their printed copy on Manning's site by creating a user account and then following instructions printed on the pBook registration insert at the front of the book. What's Inside Use group data to make individual recommendations Find logical clusters within your data Filter and refine with on-the-fly classification Free audio and video extras Table of Contents Meet Apache Mahout PART 1 RECOMMENDATIONS Introducing recommenders Representing recommender data Making recommendations Taking recommenders to production Distributing recommendation computations PART 2 CLUSTERING Introduction to clustering Representing data Clustering algorithms in Mahout Evaluating and improving clustering quality Taking clustering to production Real-world applications of clustering PART 3 CLASSIFICATION Introduction to classification Training a classifier Evaluating and tuning a classifier Deploying a classifier Case study: Shop It To Me

**React in Action** Simon and Schuster

Summary Tika in Action is a hands-on guide to content mining with Apache Tika. The book's many examples and case studies offer real-world experience from domains ranging from search engines to digital asset management and scientific data processing. About the Technology Tika is an Apache toolkit that has built into it everything you and your app need to know about file formats. Using Tika, your applications can discover and extract content from digital documents in almost any format, including exotic ones. About this Book Tika in Action is the ultimate guide to content mining using Apache Tika. You'll learn how to pull usable information from otherwise inaccessible sources, including internet media and file archives. This example-rich book teaches you to build and extend applications based on real-world experience with search engines, digital asset management, and scientific data processing. In addition to architectural overviews, you'll find detailed chapters on features like metadata extraction, automatic language detection, and custom parser development. This book is written for developers who are new to both Scala and Lift and covers just enough Scala to get you started. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Crack MS Word, PDF, HTML, and ZIP Integrate with search engines, CMS, and

other data sources Learn through experimentation Many examples This book requires no previous knowledge of Tika or text mining techniques. It assumes a working knowledge of Java.

=====

=== Table of Contents PART 1 GETTING STARTED The case for the digital Babel fish Getting started with Tika The information landscape PART 2 TIKA IN DETAIL Document type detection Content extraction Understanding metadata Language detection What's in a file? PART 3 INTEGRATION AND ADVANCED USE The big picture Tika and the Lucene search stack Extending Tika PART 4 CASE STUDIES Powering NASA science data systems Content management with Apache Jackrabbit Curating cancer research data with Tika The classic search engine example

Microservices Patterns Manning Publications

Unlock the groundbreaking advances of deep learning with this extensively revised edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In Deep Learning with Python, Second Edition you will learn: Deep learning from first principles Image classification & image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation Deep Learning with Python has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like automated language translation, image recognition, and more. Deep learning is becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach, even if you have no background in mathematics or data science. About the book Deep Learning with Python, Second Edition introduces the field of deep learning using Python and the powerful Keras library. In this new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll pick up the skills to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions Classical Mechanics Simon and Schuster Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book

gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment and administration [Play Framework Cookbook](#) O'Reilly Media Summary PostGIS in Action, Second Edition teaches readers of all levels to write spatial queries that solve real-world problems. It first gives you a background in vector-, raster-, and topology-based GIS and then quickly moves into analyzing, viewing, and mapping data. This second edition covers PostGIS 2.0 and 2.1 series, PostgreSQL 9.1, 9.2, and 9.3 features, and shows you how to integrate with other GIS tools. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book Processing data tied to location and topology requires specialized know-how. PostGIS is a free spatial database extender for PostgreSQL, every bit as good as proprietary software. With it, you can easily create location-aware queries in just a few lines of SQL code and build the back end for a mapping, raster analysis, or routing application with minimal effort. PostGIS in Action, Second Edition teaches you to solve real-world geodata problems. It first gives you a background in vector-, raster-, and topology-based GIS and then quickly moves into analyzing, viewing, and mapping data. You'll learn how to optimize queries for maximum speed, simplify geometries for greater efficiency, and create custom functions for your own applications. You'll also learn how to apply your existing GIS knowledge to PostGIS and integrate with other GIS tools. Familiarity with relational database and GIS concepts is helpful but not required. What's Inside An introduction to spatial databases Geometry, geography, raster, and topology spatial types, functions, and queries Applying PostGIS to real-world problems Extending PostGIS to web and desktop applications Updated for PostGIS 2.x and PostgreSQL 9.x About the Authors

Regina Obe and Leo Hsu are database consultants and authors. Regina is a member of the PostGIS core development team and the Project Steering Committee. Table of Contents PART 1 INTRODUCTION TO POSTGIS What is a spatial database? Spatial data types Spatial reference system considerations Working with real data Using PostGIS on the desktop Geometry and geography functions Raster functions PostGIS TIGER geocoder Geometry relationships PART 2 PUTTING POSTGIS TO WORK Proximity analysis Geometry and geography processing Raster processing Building and using topologies Organizing spatial data Query performance tuning PART 3 USING POSTGIS WITH OTHER TOOLS Extending PostGIS with pgRouting and procedural languages Using PostGIS in web applications [Apache Pulsar in Action](#) O'Reilly Media Summary Taming Text, winner of the 2013 Jolt Awards for Productivity, is a hands-on, example-driven guide to working with unstructured text in the context of real-world applications. This book explores how to automatically organize text using approaches such as full-text search, proper name recognition, clustering, tagging, information extraction, and summarization. The book guides you through examples illustrating each of these topics, as well as the foundations upon which they are built. About this Book There is so much text in our lives, we are practically drowning in it. Fortunately, there are innovative tools and techniques for managing unstructured information that can throw the smart developer a much-needed lifeline. You'll find them in this book. Taming Text is a practical, example-driven guide to working with text in real applications. This book introduces you to useful techniques like full-text search, proper name recognition, clustering, tagging, information extraction, and summarization. You'll explore real use cases as you systematically absorb the foundations upon which they are built. Written in a clear and concise style, this book avoids jargon, explaining the subject in terms you can understand without a background in statistics or natural language processing. Examples are in Java, but the concepts can be applied in any language. Written for Java developers, the book requires no prior knowledge of GWT. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. Winner of 2013 Jolt Awards: The Best Books—one of five notable books every serious programmer should read. What's Inside When to use text-taming techniques Important open-source libraries like Solr and Mahout How to build text-processing applications About the Authors Grant Ingersoll is an engineer, speaker, and trainer, a Lucene committer, and a cofounder of the Mahout machine-learning project. Thomas Morton is the primary developer of OpenNLP and Maximum Entropy. Drew Farris is a technology consultant, software developer, and contributor to Mahout, Lucene, and Solr. "Takes the mystery out of very complex processes."—From the Foreword by Liz Liddy, Dean, iSchool, Syracuse University Table of Contents Getting started taming text Foundations of taming text Searching Fuzzy string matching Identifying people, places, and things Clustering text Classification, categorization, and tagging Building an example question answering system Untamed text: exploring the next frontier [A Brain-Friendly Guide](#) Simon and Schuster Summary Enterprise OSGi in Action is a hands-on guide for developers using OSGi to build the next generation of enterprise Java applications. By presenting relevant examples and case studies, this book guides the reader through the maze of new standards and projects. About This Book Enterprise OSGi is a set of standards for building modular Java applications which integrate seamlessly with existing Java EE technologies. It extends the OSGi component framework to distributed systems.



Enterprise OSGi in Action is a hands-on guide for developers using OSGi to build enterprise Java applications. Many examples and case studies show you how to build, test, and deploy modular web applications. The book explains how to take advantage of dynamism, distribution, and automatic dependency provisioning, while still integrating with existing Java EE applications. The book is written for Java EE developers. No prior experience with OSGi is required. What's Inside Build modular applications using servlets, JSPs, WARs, and JPA Better component reuse and robustness Expert tips for Apache Aries Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Authors Holly Cummins and Tim Ward are lead engineers who regularly speak at developerWorks, DevOxx, JavaZone, and EclipseCon. Tim has written standards in both the OSGi Core and Enterprise Specifications and both authors are active Apache Aries committers. Table of Contents PART 1 PROGRAMMING BEYOND HELLO WORLD OSGi and the enterprise—why now? Developing a simple OSGi-based web application Persistence pays off Packaging your enterprise OSGi applications PART 2 BUILDING BETTER ENTERPRISE OSGI APPLICATIONS Best practices for enterprise applications Building dynamic applications with OSGi services Provisioning and resolution Tools for building and testing IDE development tools PART 3 INTEGRATING ENTERPRISE OSGI WITH EVERYTHING ELSE Hooking up remote systems with distributed OSGi Migration and integration Coping with the non-OSGi world Choosing a stack **Mahout in Action** Simon and Schuster Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on

Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance *Spring Boot in Action* Simon and Schuster Deliver lightning fast and reliable messaging for your distributed applications with the flexible and resilient Apache Pulsar platform. In *Apache Pulsar in Action* you will learn how to: Publish from Apache Pulsar into third-party data repositories and platforms Design and develop Apache Pulsar functions Perform interactive SQL queries against data stored in Apache Pulsar *Apache Pulsar in Action* is a comprehensive and practical guide to building high-traffic applications with Pulsar. You'll learn to use this mature and battle-tested platform to deliver extreme levels of speed and durability to your messaging. Apache Pulsar committer David Kjerrumgaard teaches you to apply Pulsar's seamless scalability through hands-on case studies, including IOT analytics applications and a microservices app based on Pulsar functions. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Reliable server-to-server messaging is the heart of a distributed application. Apache Pulsar is a flexible real-time messaging platform built to run on Kubernetes and deliver the scalability and resilience required for cloud-based systems. Pulsar supports both streaming and message queuing, and unlike other solutions, it can communicate over multiple protocols including MQTT, AMQP, and Kafka's binary protocol. About the book *Apache Pulsar in Action* teaches you to build scalable streaming messaging systems using Pulsar. You'll start with a rapid introduction to enterprise messaging and discover the unique benefits of Pulsar. Following crystal-clear explanations and engaging examples, you'll use the Pulsar Functions framework to develop a microservices-based application. Real-world case studies illustrate how to implement the most important messaging design patterns. What's inside Publish from Pulsar into third-party data repositories and platforms Design and develop Apache Pulsar functions Create an event-driven food delivery application About the reader Written for experienced Java developers. No prior knowledge of Pulsar required. About the author David Kjerrumgaard is a committer on the Apache Pulsar project. He currently serves as a Developer Advocate for StreamNative, where he develops Pulsar best practices and solutions. Table of Contents PART 1 GETTING STARTED WITH APACHE PULSAR 1 Introduction to Apache Pulsar 2 Pulsar concepts and architecture 3 Interacting with Pulsar PART 2 APACHE PULSAR DEVELOPMENT ESSENTIALS 4 Pulsar functions 5 Pulsar IO connectors 6 Pulsar security 7 Schema registry PART 3 HANDS-ON APPLICATION DEVELOPMENT WITH APACHE PULSAR 8 Pulsar Functions patterns 9 Resiliency patterns 10 Data access 11 Machine learning in Pulsar 12 Edge analytics *CMIS and Apache Chemistry in Action* Manning Publications Introduces the build tool for Java application development, covering both user defined and built-in tasks. *Covers MongoDB version 3.0* Simon and Schuster Summary HBase in Action has all the knowledge you need to design, build, and run applications using HBase. First, it introduces you to the fundamentals of distributed systems and large scale data handling. Then, you'll explore real-world applications and code samples with just enough theory to

understand the practical techniques. You'll see how to build applications with HBase and take advantage of the MapReduce processing framework. And along the way you'll learn patterns and best practices. About the Technology HBase is a NoSQL storage system designed for fast, random access to large volumes of data. It runs on commodity hardware and scales smoothly from modest datasets to billions of rows and millions of columns. About this Book HBase in Action is an experience-driven guide that shows you how to design, build, and run applications using HBase. First, it introduces you to the fundamentals of handling big data. Then, you'll explore HBase with the help of real applications and code samples and with just enough theory to back up the practical techniques. You'll take advantage of the MapReduce processing framework and benefit from seeing HBase best practices in action. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside When and how to use HBase Practical examples Design patterns for scalable data systems Deployment, integration, and design Written for developers and architects familiar with data storage and processing. No prior knowledge of HBase, Hadoop, or MapReduce is required. Table of Contents PART 1 HBASE FUNDAMENTALS Introducing HBase Getting started Distributed HBase, HDFS, and MapReduce PART 2 ADVANCED CONCEPTS HBase table design Extending HBase with coprocessors Alternative HBase clients PART 3 EXAMPLE APPLICATIONS HBase by example: OpenTSDB Scaling GIS on HBase PART 4 OPERATIONALIZING HBASE Deploying HBase Operations

[Kubernetes in Action](#) Simon and Schuster

Upgrade your machine learning models with graph-based algorithms, the perfect structure for complex and interlinked data. Summary In Graph-Powered Machine Learning, you will learn: The lifecycle of a machine learning project Graphs in big data platforms Data source modeling using graphs Graph-based natural language processing, recommendations, and fraud detection techniques Graph algorithms Working with Neo4J Graph-Powered Machine Learning teaches to use graph-based algorithms and data organization strategies to develop superior machine learning applications. You'll dive into the role of graphs in machine learning and big data platforms, and take an in-depth look at data source modeling, algorithm design, recommendations, and fraud detection. Explore end-to-end projects that illustrate architectures and help you optimize with best design practices. Author Alessandro Negro's extensive experience shines through in every chapter, as you learn from examples and concrete scenarios based on his work with real clients! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Identifying relationships is the foundation of machine learning. By recognizing and analyzing the connections in your data, graph-centric algorithms like K-nearest neighbor or PageRank radically improve the effectiveness of ML applications. Graph-based machine learning techniques offer a powerful new perspective for machine learning in social networking, fraud detection, natural language processing, and recommendation systems. About the book Graph-Powered Machine Learning teaches you how to exploit the natural relationships in structured and unstructured datasets using graph-oriented machine learning algorithms and tools. In this authoritative book, you'll master the architectures and design practices of graphs, and avoid common pitfalls. Author Alessandro Negro explores examples from real-world applications that connect GraphML concepts to real world tasks. What's inside Graphs in big data platforms Recommendations, natural language processing, fraud detection Graph algorithms Working with the Neo4J graph database About

the reader For readers comfortable with machine learning basics. About the author Alessandro Negro is Chief Scientist at GraphAware. He has been a speaker at many conferences, and holds a PhD in Computer Science. Table of Contents PART 1 INTRODUCTION 1 Machine learning and graphs: An introduction 2 Graph data engineering 3 Graphs in machine learning applications PART 2 RECOMMENDATIONS 4 Content-based recommendations 5 Collaborative filtering 6 Session-based recommendations 7 Context-aware and hybrid recommendations PART 3 FIGHTING FRAUD 8 Basic approaches to graph-powered fraud detection 9 Proximity-based algorithms 10 Social network analysis against fraud PART 4 TAMING TEXT WITH GRAPHS 11 Graph-based natural language processing 12 Knowledge graphs [Relevant Search](#) Simon and Schuster

Users expect search to be simple: They enter a few terms and expect perfectly-organized, relevant results instantly. But behind this simple user experience, complex machinery is at work. Whether using Elasticsearch, Solr, or another search technology, the solution is never one size fits all. Returning the right search results requires conveying domain knowledge and business rules in the search engine's data structures, text analytics, and results ranking capabilities. Relevant Search demystifies relevance work. Using Elasticsearch, it tells how to return engaging search results to users, helping readers understand and leverage the internals of Lucene-based search engines. The book walks through several real-world problems using a cohesive philosophy that combines text analysis, query building, and score shaping to express business ranking rules to the search engine. It outlines how to guide the engineering process by monitoring search user behavior and shifting the enterprise to a search-first culture focused on humans, not computers. It also shows how the search engine provides a deeply pluggable platform for integrating search ranking with machine learning, ontologies, personalization, domain-specific expertise, and other enriching sources. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

[Spring Roo in Action](#) Simon and Schuster

This book is in Packt's Cookbook series. A Packt Cookbook contains recipes for solutions to the most important problems you face when working with a topic. Inside the Cookbook you will find: A straightforward and easy to follow format, A selection of the most important tasks and problems, Carefully organized instructions for solving the problem efficiently, Clear explanations of what you did, Details for applying the solution to other situations. This is the ideal book for people who have already written a first application with the Play Framework or have just finished reading through the documentation. In other words - anyone who is ready to get to grips with Play! Having a basic knowledge of Java is good, as well as well as some web developer skills - HTML and JavaScript

[Enterprise OSGi In Action](#) Manning Publications

Summary A developer-focused guide to writing applications using Spring Boot. You'll learn how to bypass the tedious configuration steps so that you can concentrate on your application's behavior. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The Spring Framework simplifies enterprise Java development, but it does require lots of tedious configuration work. Spring Boot radically streamlines spinning up a Spring application. You get automatic configuration and a model with established conventions for build-time and runtime dependencies. You also get a handy command-line interface you can use to write scripts in Groovy. Developers who use Spring Boot often say that they can't imagine going back to hand configuring their applications. About the Book Spring Boot in

Action is a developer-focused guide to writing applications using Spring Boot. In it, you'll learn how to bypass configuration steps so you can focus on your application's behavior. Spring expert Craig Walls uses interesting and practical examples to teach you both how to use the default settings effectively and how to override and customize Spring Boot for your unique environment. Along the way, you'll pick up insights from Craig's years of Spring development experience. What's Inside Develop Spring apps more efficiently Minimal to no configuration Runtime metrics with the Actuator Covers Spring Boot 1.3 About the Reader Written for readers familiar with the Spring Framework. About the Author Craig Walls is a software developer, author of the popular book Spring in Action, Fourth Edition, and a frequent speaker at conferences. Table of Contents Bootstarting Spring Developing your first Spring Boot application Customizing configuration Testing with Spring Boot Getting Groovy with the Spring Boot CLI Applying Grails in Spring Boot Taking a peek inside with the Actuator Deploying Spring Boot applications APPENDIXES Spring Boot developer tools Spring Boot starters Configuration properties Spring Boot dependencies

*A Comprehensive Guide to Building Real-World NLP Systems*  
Simon and Schuster

Summary Solr in Action is a comprehensive guide to implementing scalable search using Apache Solr. This clearly written book walks you through well-documented examples ranging from basic keyword searching to scaling a system for billions of documents and queries. It will give you a deep understanding of how to implement core Solr capabilities. About the Book Whether you're handling big (or small) data, managing documents, or building a website, it is important to be able to quickly search through your content and discover meaning in it. Apache Solr is your tool: a ready-to-deploy, Lucene-based, open source, full-text search engine. Solr can scale across many servers to enable real-time queries and data analytics across billions of documents. Solr in Action teaches you to implement scalable search using Apache Solr. This easy-to-read guide balances conceptual discussions with practical examples to show you how to implement all of Solr's core capabilities. You'll master topics like text analysis, faceted search, hit highlighting, result grouping, query suggestions, multilingual search, advanced geospatial and data operations, and relevancy tuning. This book assumes basic knowledge of Java and standard database technology. No prior knowledge of Solr or Lucene is required. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside How to scale Solr for big data Rich real-world examples Solr as a NoSQL data store Advanced multilingual, data, and relevancy tricks Coverage of versions through Solr 4.7 About the Authors Trey Grainger is a director of engineering at CareerBuilder. Timothy Potter is a senior member of the engineering team at LucidWorks. The authors work on the scalability and reliability of Solr, as well as on recommendation engine and big data analytics technologies. Table of Contents PART 1 MEET SOLR Introduction to Solr Getting to know Solr Key Solr concepts Configuring Solr Indexing Text analysis PART 2 CORE SOLR CAPABILITIES Performing queries and handling results Faceted search Hit highlighting Query suggestions Result grouping/field collapsing Taking Solr to production PART 3 TAKING SOLR TO THE NEXT LEVEL SolrCloud Multilingual search Complex query operations Mastering relevancy

Grokking Machine Learning Solr in Action

Summary CMIS and Apache Chemistry in Action is a comprehensive guide to the CMIS standard and related ECM concepts, written by the authors of the standard. In it, you'll tackle hands-on examples for building applications on CMIS

repositories from both the client and the server sides. You'll learn how to create new content-centric applications that install and run in any CMIS-compliant repository. About The Technology Content Management Interoperability Services (CMIS) is an OASIS standard for accessing content management systems. It specifies a vendor-and language-neutral way to interact with any compliant content repository. Apache Chemistry provides complete reference implementations of the CMIS standard with robust APIs for developers writing tools, applications, and servers. About This Book CMIS and Apache Chemistry in Action is a comprehensive guide to the CMIS standard and related ECM concepts. In it, you'll find clear teaching and instantly useful examples for building content-centric client and server-side applications that run against any CMIS-compliant repository. In fact, using the CMIS Workbench and the InMemory Repository from Apache Chemistry, you'll have running code talking to a real CMIS server by the end of chapter 1. This book requires some familiarity with content management systems and a standard programming language like Java or C#. No exposure to CMIS or Apache Chemistry is assumed. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside The only CMIS book endorsed by OASIS Complete coverage of the CMIS 1.0 and 1.1 specifications Cookbook-style tutorials and real-world examples About the Authors Florian Müller, Jay Brown, and Jeff Potts are among the original authors, contributors, and leaders of Apache Chemistry and the OASIS CMIS specification. They continue to shape CMIS implementations at Alfresco, IBM, and SAP. Table of Contents PART 1 UNDERSTANDING CMIS Introducing CMIS Exploring the CMIS domain model Creating, updating, and deleting objects with CMIS CMIS metadata: types and properties Query PART 2 HANDS-ON CMIS CLIENT DEVELOPMENT Meet your new project: The Blend The Blend: read and query functionality The Blend: create, update, and delete functionality Using other client libraries Building mobile apps with CMIS PART 3 ADVANCED TOPICS CMIS bindings Security and control Performance Building a CMIS server

**Covers Apache Spark 3 with Examples in Java, Python, and Scala** Tata McGraw-Hill Education

Summary Elasticsearch in Action teaches you how to build scalable search applications using Elasticsearch. You'll ramp up fast, with an informative overview and an engaging introductory example. Within the first few chapters, you'll pick up the core concepts you need to implement basic searches and efficient indexing. With the fundamentals well in hand, you'll go on to gain an organized view of how to optimize your design. Perfect for developers and administrators building and managing search-oriented applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Modern search seems like magic—you type a few words and the search engine appears to know what you want. With the Elasticsearch real-time search and analytics engine, you can give your users this magical experience without having to do complex low-level programming or understand advanced data science algorithms. You just install it, tweak it, and get on with your work. About the Book Elasticsearch in Action teaches you how to write applications that deliver professional quality search. As you read, you'll learn to add basic search features to any application, enhance search results with predictive analysis and relevancy ranking, and use saved data from prior searches to give users a custom experience. This practical book focuses on Elasticsearch's REST API via HTTP. Code snippets are written mostly in bash using cURL, so they're easily translatable to other languages. What's Inside What is a great search application? Building scalable search solutions Using



Elasticsearch with any language Configuration and tuning About the Reader For developers and administrators building and managing search-oriented applications. About the Authors Radu Gheorghe is a search consultant and software engineer. Matthew Lee Hinman develops highly available, cloud-based systems. Roy Russo is a specialist in predictive analytics. Table of Contents PART 1 CORE ELASTICSEARCH FUNCTIONALITY Introducing

Elasticsearch Diving into the functionality Indexing, updating, and deleting data Searching your data Analyzing your data Searching with relevancy Exploring your data with aggregations Relations among documents PART 2 ADVANCED ELASTICSEARCH FUNCTIONALITY Scaling out Improving performance Administering your cluster