
Online Library Mongo Db Ieee Paper 201

When people should go to the books stores, search creation by shop, shelf by shelf, it is essentially problematic. This is why we provide the ebook compilations in this website. It will extremely ease you to look guide **Mongo Db Ieee Paper 201** as you such as.

By searching the title, publisher, or authors of guide you in reality want, you can discover them rapidly. In the house, workplace, or perhaps in your method can be all best place within net connections. If you aspire to download and install the Mongo Db Ieee Paper 201, it is unconditionally easy then, past currently we extend the join to buy and make bargains to download and install Mongo Db Ieee Paper 201 as a result simple!

HPTTJC - ANAYA CARLEE

This book presents recent advances on IoT and connected technologies. We are currently in the midst of the Fourth Industrial Revolution, and IoT is having the most significant impact on our society. The recent adoption of a variety of enabling wireless communication technologies like RFID tags, BLE, ZigBee, etc., embedded sensor and actuator nodes, and various protocols like CoAP, MQTT, DNS, etc., has made the Internet of things (IoT) step out of its infancy. Internet of things (IoT) and connecting technologies are already having profound effects on the different parts of society like the government, health care, businesses, and personal lives. 6th International Conference on Internet of Things and Connected Technologies (ICIoTCT), 2021, was a platform to discuss and feature research on topics such as augmented reality, sensor networks, and wearable technology. This book is ideally designed for marketing managers, business professionals, researchers, academicians, and graduate-level students seeking to learn how IoT and connecting technologies increase the amount of data gained through devices, enhance customer experience, and widen the scope of IoT analytics in enhancing customer marketing outcomes.

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

This book includes a selection of articles from The 2019 World Conference on Information Systems and Technologies (WorldCIST'19), held from April 16 to 19, at La Toja, Spain. WorldCIST is a global forum for researchers and practitioners to present and discuss recent results and innovations, current trends, professional experiences and challenges in modern information systems and technologies research, together with their technological development and applications. The book covers a number of topics, including A) Information and Knowledge Management; B) Organizational Models and Information Systems; C) Software and Systems Modeling; D) Software Systems, Architectures, Applications and Tools; E) Multimedia Systems and Applications; F) Computer Networks, Mobility and Pervasive Systems; G) Intelligent and Decision Support Systems; H) Big Data Analytics and Applications; I) Human-Computer Interaction; J) Ethics, Computers & Security; K) Health Informatics; L) Information Technologies in Education; M) Information Technologies in Radiocommunications; and N) Technologies for Biomedical Applications.

Abstract This IBM® Redbooks® publication presents a general introduction to the latest IBM tape and tape library technologies. Featured tape technologies include the IBM LTO Ultrium and Enterprise 3592 tape drives, and their implementation in IBM tape libraries. This 16th edition introduces the new TS1160 tape drive with up to 20 TB capacity on JE media and the latest updates to the IBM TS4500 and TS4300 tape libraries. It includes generalized sections about Small Computer System Interface (SCSI) and Fibre Channel connections, and multipath architecture configurations. This book also covers tools and techniques for library management. It is intended for anyone who wants to understand more about IBM tape products and their implementation. It is suitable for IBM clients, IBM Business Partners, IBM specialist sales representatives, and technical specialists. If you do not have a background in computer tape storage products, you might need to read other sources of information. In the interest of being concise, topics that are generally understood are not covered in detail.

Cyber-security is a matter of rapidly growing importance in industry and government. This book provides insight into a range of data science techniques for addressing these pressing concerns. The application of statistical and broader data science techniques provides an exciting growth area in the design of cyber defences. Networks of connected devices, such as enterprise computer networks or the wider so-called Internet of Things, are all vulnerable to misuse and attack, and data science methods offer the promise to detect such behaviours from the vast collections of cyber traffic data sources that can be obtained. In many cases, this is achieved through anomaly detection of unusual behaviour against understood statistical models of normality. This volume presents contributed papers from an international conference of the same name held at Imperial College. Experts from the

field have provided their latest discoveries and review state of the art technologies.

The topic of NoSQL databases has recently emerged, to face the Big Data challenge, namely the ever increasing volume of data to be handled. It is now recognized that relational databases are not appropriate in this context, implying that new database models and techniques are needed. This book presents recent research works, covering the following basic aspects: semantic data management, graph databases, and big data management in cloud environments. The chapters in this book report on research about the evolution of basic concepts such as data models, query languages, and new challenges regarding implementation issues.

Summary Getting MEAN, Second Edition teaches you how to develop full-stack web applications using the MEAN stack. This edition was completely revised and updated to cover MongoDB 4, Express 4, Angular 7, Node 11, and the latest mainstream release of JavaScript ES2015. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Juggling languages mid-application can radically slow down a full-stack web project. The MEAN stack—MongoDB, Express, Angular, and Node—uses JavaScript end to end, maximizing developer productivity and minimizing context switching. And you'll love the results! MEAN apps are fast, powerful, and beautiful. About the Book Getting MEAN, Second Edition teaches you how to develop full-stack web applications using the MEAN stack. Practical from the very beginning, the book helps you create a static site in Express and Node. Expanding on that solid foundation, you'll integrate a MongoDB database, build an API, and add an authentication system. Along the way, you'll get countless pro tips for building dynamic and responsive data-driven web applications! What's inside MongoDB 4, Express 4, Angular 7, and Node.js 11 MEAN stack architecture Mobile-ready web apps Best practices for efficiency and reusability About the Reader Readers should be comfortable with standard web application designs and ES2015-style JavaScript. About the Author Simon Holmes and Clive Harber are full-stack developers with decades of experience in JavaScript and other leading-edge web technologies. Table of Contents PART 1 - SETTING THE BASELINE Introducing full-stack development Designing a MEAN stack architecture PART 2 - BUILDING A NODE WEB APPLICATION Creating and setting up a MEAN project Building a static site with Node and Express Building a data model with MongoDB and Mongoose Writing a REST API: Exposing the MongoDB database to the application Consuming a REST API: Using an API from inside Express PART 3 - ADDING A DYNAMIC FRONT END WITH ANGULAR Creating an Angular application with TypeScript Building a single-page application with Angular: Foundations Building a single-page application with Angular: The next level PART 4 - MANAGING AUTHENTICATION AND USER SESSIONS Authenticating users, managing sessions, and securing APIs Using an authentication API in Angular applications

The book aims to provide a broad overview of various topics of the Internet of Things (IoT) from the research and development priorities to enabling technologies, architecture, security, privacy, interoperability and industrial applications. It is intended to be a stand-alone book in a series that covers the Internet of Things activities of the IERC - Internet of Things European Research Cluster - from technology to international cooperation and the global "state of play." The book builds on the ideas put forward by the European Research Cluster on the Internet of Things Strategic Research and Innovation Agenda and presents views and state of the art results on the challenges facing the research, development and deployment of IoT at the global level. Today we see the integration of Industrial, Business and Consumer Internet which is bringing together the Internet of People, Internet of Things, Internet of Energy, Internet of Vehicles, Internet of Media, Services and Enterprises in forming the backbone of the digital economy, the digital society and the foundation for the future knowledge and innovation based economy. These developments are supporting solutions for the emerging challenges of public health, aging population, environmental protection and climate change, the conservation of energy and scarce materials, enhancements to safety and security and the continuation and growth of economic prosperity. Penetration of smartphones and advances in nanoelectronics, cyber-physical systems, wireless communication, software, and Cloud computing technology will be the main drivers for IoT development. The IoT contribution is seen in the increased value of information created by the number of interconnections among things and the transformation of the processed information into knowledge shared into the Internet of Everything. The connected devices are part of ecosystems connecting people, processes, data, and things which are communicating in the Cloud using the increased storage and computing power while attempting to standardize communication and metadata. In this context, the next generation of Cloud computing technologies will need to be flexible enough to scale autonomously, adaptive enough to handle constantly changing connections and resilient enough to stand up to the huge flows of data that will occur. In 2025, analysts forecast that there will be six devices per human on the planet, which means around 50 billion more connected devices over the next 12 years. The Internet of Things market is connected to this anticipated device growth from industrial Machine to Machine (M2M) systems, smart meters and wireless sensors. Internet of Things technology will generate new services and new interfaces by creating smart environments and smart spaces with applications ranging from Smart Cities, Smart Transport, Buildings, Energy, Grid, to Smart Health and Life.

Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor

workloads, tune queries, and connect more applications; learn how other organizations apply Trino

This Springer Brief provides a comprehensive overview of the background and recent developments of big data. The value chain of big data is divided into four phases: data generation, data acquisition, data storage and data analysis. For each phase, the book introduces the general background, discusses technical challenges and reviews the latest advances. Technologies under discussion include cloud computing, Internet of Things, data centers, Hadoop and more. The authors also explore several representative applications of big data such as enterprise management, online social networks, healthcare and medical applications, collective intelligence and smart grids. This book concludes with a thoughtful discussion of possible research directions and development trends in the field. Big Data: Related Technologies, Challenges and Future Prospects is a concise yet thorough examination of this exciting area. It is designed for researchers and professionals interested in big data or related research. Advanced-level students in computer science and electrical engineering will also find this book useful.

This book constitutes revised selected papers from the 26th Argentine Congress on Computer Science, CACIC 2020, held in San Justo, Buenos Aires, Argentina in October 2020. Due to the COVID-19 pandemic the conference was held in a virtual mode. The 21 full papers and 3 short papers presented in this volume were carefully reviewed and selected from a total of 118 submissions. They were organized in topical sections named: intelligent agents and systems; distributed and parallel processing; computer technology applied to education; graphic computation, images and visualization; software engineering; databases and data mining; hardware architectures, networks, and operating systems; innovation in software systems; signal processing and real-time systems; innovation in computer science education; computer security; and digital governance and smart cities.

Dr. Greg Zacharias, former Chief Scientist of the United States Air Force (2015-18), explores next steps in autonomous systems (AS) development, fielding, and training. Rapid advances in AS development and artificial intelligence (AI) research will change how we think about machines, whether they are individual vehicle platforms or networked enterprises. The payoff will be considerable, affording the US military significant protection for aviators, greater effectiveness in employment, and unlimited opportunities for novel and disruptive concepts of operations. Autonomous Horizons: The Way Forward identifies issues and makes recommendations for the Air Force to take full advantage of this transformational technology.

Internet-of-Things (IoT) Analytics are an integral element of most IoT applications, as it provides the means to extract knowledge, drive actuation services and optimize decision making. IoT analytics will be a major contributor to IoT business value in the coming years, as it will enable organizations to process and fully leverage large amounts of IoT data, which are nowadays largely underutilized. The Building Blocks of IoT Analytics is devoted to the presentation the main technology building blocks that comprise advanced IoT analytics systems. It introduces IoT analytics as a special case of BigData analytics and accordingly presents leading edge technologies that can be deployed in order to successfully confront the main challenges of IoT analytics applications. Special emphasis is paid in the presentation of technologies for IoT streaming and semantic interoperability across diverse IoT streams. Furthermore, the role of cloud computing and BigData technologies in IoT analytics are presented, along with practical tools for implementing, deploying and operating non-trivial IoT applications. Along with the main building blocks of IoT analytics systems and applications, the book presents a series of practical applications, which illustrate the use of these technologies in the scope of pragmatic applications. Technical topics discussed in the book include: Cloud Computing and BigData for IoT analytics Searching the Internet of Things Development Tools for IoT Analytics Applications IoT Analytics-as-a-Service Semantic Modelling and Reasoning for IoT Analytics IoT analytics for Smart Buildings IoT analytics for Smart Cities Operationalization of IoT analytics Ethical aspects of IoT analytics This book contains both research oriented and applied articles on IoT analytics, including several articles reflecting work undertaken in the scope of recent European Commission funded projects in the scope of the FP7 and H2020 programmes. These articles present results of these projects on IoT analytics platforms and applications. Even though several articles have been contributed by different authors, they are structured in a well thought order that facilitates the reader either to follow the evolution of the book or to focus on specific topics depending on his/her background and interest in IoT and IoT analytics technologies. The compilation of these articles in this edited volume has been largely motivated by the close collaboration of the co-authors in the scope of working groups and IoT events organized by the Internet-of-Things Research Cluster (IERC), which is currently a part of EU's Alliance for Internet of Things Innovation (AIOTI).

Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a *nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.

This book gathers selected research papers presented at the International Conference on Recent Trends in Machine Learning, IOT, Smart Cities & Applications (ICMISC 2020), held on 29-30 March 2020 at CMR Institute of Technology, Hyderabad, Telangana, India. Discussing current trends in machine learning, Internet of things, and smart cities applications, with a focus on multi-disciplinary research in the area of artificial intelligence and cyber-physical systems, this book is a valuable resource for scientists, research scholars and PG students wanting formulate their research ideas and find the future directions in these areas. Further, it serves as a reference work anyone wishing to understand the latest technologies used by practic-

ing engineers around the globe.

Design, administer, and deploy high-volume and fault-tolerant database applications using MongoDB 4.x Key Features Build a powerful and scalable MongoDB database using real industry data Understand the process of designing NoSQL schema with the latest release of MongoDB 4.x Explore the ins and outs of MongoDB, including queries, replication, sharding, and vital admin tasks Book Description When it comes to managing a high volume of unstructured and non-relational datasets, MongoDB is the defacto database management system (DBMS) for DBAs and data architects. This updated book includes the latest release and covers every feature in MongoDB 4.x, while helping you get hands-on with building a MongoDB database app. You'll get to grips with MongoDB 4.x concepts such as indexes, database design, data modeling, authentication, and aggregation. As you progress, you'll cover tasks such as performing routine operations when developing a dynamic database-driven website. Using examples, you'll learn how to work with queries and regular database operations. The book will not only guide you through design and implementation, but also help you monitor operations to achieve optimal performance and secure your MongoDB database systems. You'll also be introduced to advanced techniques such as aggregation, map-reduce, complex queries, and generating ad hoc financial reports on the fly. Later, the book shows you how to work with multiple collections as well as embedded arrays and documents, before finally exploring key topics such as replication, sharding, and security using practical examples. By the end of this book, you'll be well-versed with MongoDB 4.x and be able to perform development and administrative tasks associated with this NoSQL database. What you will learn Understand how to configure and install MongoDB 4.x Build a database-driven website using MongoDB as the backend Perform basic database operations and handle complex MongoDB queries Develop a successful MongoDB database design for large corporate customers with complex requirements Secure MongoDB database systems by establishing role-based access control with X.509 transport-level security Optimize reads and writes directed to a replica set or sharded cluster Perform essential MongoDB administration tasks Maintain database performance through monitoring Who this book is for This book is a MongoDB tutorial for DevOps engineers, database developers, database administrators, system administrators and those who are just getting started with NoSQL and looking to build document-oriented databases and gain real-world experience in managing databases using MongoDB. Basic knowledge of databases and Python is required to get started with this DBMS book.

This two-volume set (CCIS 1075 and CCIS 1076) constitutes the refereed proceedings of the Third International Conference on Advanced Informatics for Computing Research, ICAICR 2019, held in Shimla, India, in June 2019. The 78 revised full papers presented were carefully reviewed and selected from 382 submissions. The papers are organized in topical sections on computing methodologies; hardware; information systems; networks; software and its engineering.

This book comprehensively reviews searchable encryption, which represents a series of research developments that directly enable search functionality over encrypted data. The book majorly covers: 1) the design and implementation of encrypted search algorithms, data structures, and systems that facilitate various forms of search over always-encrypted databases; 2) different threat models, assumptions, and the related security guarantees, when using searchable encryption in the real-world settings; and 3) latest efforts in building full-fledged encrypted database systems that draw insights from searchable encryption constructions. The book fits in the timely context, where the necessity of safeguarding important and sensitive data has been globally recognized. Traditional security measures, such as storing data behind network firewalls and layers of access control mechanisms to keep attackers out, are no longer sufficient to cope with the expanding landscape of surging cyber threats. There is an urgent call to keep sensitive data always encrypted to protect the data at rest, in transit, and in use. Doing so guarantees data confidentiality for owners, even if the data is out of their hands, e.g., hosted at in-the-cloud databases. The daunting challenge is how to perform computation over encrypted data. As we unfold in this book, searchable encryption, as a specific line of research in this broadly defined area, has received tremendous advancements over the past decades. This book is majorly oriented toward senior undergraduates, graduate students, and researchers, who want to work in the field and need extensive coverage of encrypted database research. It also targets security practitioners who want to make well-informed deployment choices of the latest advancements in searchable encryption for their targeted applications. Hopefully, this book will be beneficial in both regards.

There are many excellent R resources for visualization, data science, and package development. Hundreds of scattered vignettes, web pages, and forums explain how to use R in particular domains. But little has been written on how to simply make R work effectively—until now. This hands-on book teaches novices and experienced R users how to write efficient R code. Drawing on years of experience teaching R courses, authors Colin Gillespie and Robin Lovelace provide practical advice on a range of topics—from optimizing the set-up of RStudio to leveraging C++—that make this book a useful addition to any R user's bookshelf. Academics, business users, and programmers from a wide range of backgrounds stand to benefit from the guidance in Efficient R Programming. Get advice for setting up an R programming environment Explore general programming concepts and R coding techniques Understand the ingredients of an efficient R workflow Learn how to efficiently read and write data in R Dive into data carpentry—the vital skill for cleaning raw data Optimize your code with profiling, standard tricks, and other methods Determine your hardware capabilities for handling R computation Maximize the benefits of collaborative R programming Accelerate your transition from R hacker to R programmer

This open access book presents the first comprehensive overview of general methods in Automated Machine Learning (AutoML), collects descriptions of existing systems based on these methods, and discusses the first series of international challenges of AutoML systems. The recent success of commercial ML applications and the rapid growth of the field has created a high demand for off-the-shelf ML methods that can be used easily and without expert knowledge. However, many of the recent machine learning successes crucially rely on human experts, who manually select appropriate ML architectures (deep learning architectures or more traditional ML workflows) and their hyperparameters. To overcome this problem, the field of AutoML targets a progressive automation of machine learning, based on principles from optimization and machine learning itself. This book serves as a point of entry into this quickly-developing field for researchers and advanced students alike, as well as providing a reference for practitioners aiming to use AutoML in their work.

The rise of the internet and mobile telecommunications has created the possibility of using large datasets to understand behavior at unprecedented levels of temporal and geographic resolution. Online social networks attract the most users, though users of these new technologies provide their data through multiple sources, e.g. call detail records, blog posts, web forums, and content aggregation sites. These data allow scholars to adjudicate be-

tween competing theories as well as develop new ones, much as the microscope facilitated the development of the germ theory of disease. Of those networks, Twitter presents an ideal combination of size, international reach, and data accessibility that make it the preferred platform in academic studies. Acquiring, cleaning, and analyzing these data, however, require new tools and processes. This Element introduces these methods to social scientists and provides scripts and examples for downloading, processing, and analyzing Twitter data.

This book discusses the advanced databases for the cloud-based application known as NoSQL. It will explore the recent advancements in NoSQL database technology. Chapters on structured, unstructured and hybrid databases will be included to explore bigdata analytics, bigdata storage and processing. The book is likely to cover a wide range of topics such as cloud computing, social computing, bigdata and advanced databases processing techniques.

This two-volume book presents outcomes of the 7th International Conference on Soft Computing for Problem Solving, SocProS 2017. This conference is a joint technical collaboration between the Soft Computing Research Society, Liverpool Hope University (UK), the Indian Institute of Technology Roorkee, the South Asian University New Delhi and the National Institute of Technology Silchar, and brings together researchers, engineers and practitioners to discuss thought-provoking developments and challenges in order to select potential future directions. The book presents the latest advances and innovations in the interdisciplinary areas of soft computing, including original research papers in the areas including, but not limited to, algorithms (artificial immune systems, artificial neural networks, genetic algorithms, genetic programming, and particle swarm optimization) and applications (control systems, data mining and clustering, finance, weather forecasting, game theory, business and forecasting applications). It is a valuable resource for both young and experienced researchers dealing with complex and intricate real-world problems for which finding a solution by traditional methods is a difficult task.

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational “NoSQL” databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage’s pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

Introductory, theory-practice balanced text teaching the fundamentals of databases to advanced undergraduates or graduate students in information systems or computer science.

The three-volume set constitutes the proceedings of the 16th International Conference on Wireless Algorithms, Systems, and Applications, WASA 2021, which was held during June 25-27, 2021. The conference took place in Nanjing, China. The 103 full and 57 short papers presented in these proceedings were carefully reviewed and selected from 315 submissions. The contributions in Part II of the set are subdivided into the following topical sections: Scheduling & Optimization II; Security; Data Center Networks and Cloud Computing; Privacy-Aware Computing; Internet of Vehicles; Visual Computing for IoT; Mobile Ad-Hoc Networks.

Congratulations! You completed the MongoDB application within the given tight timeframe and there is a party to celebrate your application’s release into production. Although people are congratulating you at the celebration, you are feeling some uneasiness inside. To complete the project on time required making a lot of assumptions about the data, such as what terms meant and how calculations are derived. In addition, the poor documentation about the application will be of limited use to the support team, and not investigating all of the inherent rules in the data may eventually lead to poorly-performing structures in the not-so-distant future. Now, what if you had a time machine and could go back and read this book. You would learn that even NoSQL databases like MongoDB require some level of data modeling. Data modeling is the process of learning about the data, and regardless of technology, this process must be performed for a successful application. You would learn the value of conceptual, logical, and physical data modeling and how each stage increases our knowledge of the data and reduces assumptions and poor design decisions. Read this book to learn how to do data modeling for MongoDB applications, and accomplish these five objectives: Understand how data modeling contributes to the process of learning about the data, and is, therefore, a required technique, even when the resulting database is not relational. That is, NoSQL does not mean No-DataModeling! Know how NoSQL databases differ from traditional relational databases, and where MongoDB fits. Explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts, and learn the basics of adding, querying, updating, and deleting data in MongoDB. Practice a streamlined, template-driven approach to performing conceptual, logical, and physical data modeling. Recognize that data modeling does not always have to lead to traditional data models! Distinguish top-down from bottom-up development approaches and complete a top-down case study which ties all of the modeling techniques together. This book is written for anyone who is working with, or will be working with MongoDB, including business analysts, data modelers, database administrators, developers, project managers, and data scientists. There are three sections: In Section I, Getting Started, we will reveal the power of data modeling and the tight connections to data models that exist when designing any type of database (Chapter 1), compare NoSQL with traditional relational databases and where MongoDB fits (Chapter 2), explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts (Chapter 3), and explain the basics of adding, querying, updating, and deleting data in MongoDB (Chapter 4). In Section II, Levels of Granularity, we cover Conceptual Data Modeling (Chapter 5), Logical Data Modeling (Chapter 6), and Physical Data Modeling (Chapter 7). Notice the “ing” at the end of each of these chapters. We focus on the process of building each of these models, which is where we gain essential business knowledge. In Section III, Case Study, we will explain both top down and bottom up development approaches and go through a top down case study where we start with business re-

quirements and end with the MongoDB database. This case study will tie together all of the techniques in the previous seven chapters. Nike Senior Data Architect Ryan Smith wrote the foreword. Key points are included at the end of each chapter as a way to reinforce concepts. In addition, this book is loaded with hands-on exercises, along with their answers provided in Appendix A. Appendix B contains all of the book’s references and Appendix C contains a glossary of the terms used throughout the text.

This book is composed of a selection of articles from The 2021 World Conference on Information Systems and Technologies (WorldCIST’21), held online between 30 and 31 of March and 1 and 2 of April 2021 at Hangra de Heroismo, Terceira Island, Azores, Portugal. WorldCIST is a global forum for researchers and practitioners to present and discuss recent results and innovations, current trends, professional experiences and challenges of modern information systems and technologies research, together with their technological development and applications. The main topics covered are: A) Information and Knowledge Management; B) Organizational Models and Information Systems; C) Software and Systems Modeling; D) Software Systems, Architectures, Applications and Tools; E) Multimedia Systems and Applications; F) Computer Networks, Mobility and Pervasive Systems; G) Intelligent and Decision Support Systems; H) Big Data Analytics and Applications; I) Human-Computer Interaction; J) Ethics, Computers & Security; K) Health Informatics; L) Information Technologies in Education; M) Information Technologies in Radiocommunications; N) Technologies for Biomedical Applications.

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.

This volume constitutes refereed proceedings of the Third International Conference on Smart Applications and Data Analysis, SADASC 2020, held in Marrakesh, Morocco. Due to the COVID-19 pandemic the conference has been postponed to June 2020. The 24 full papers and 3 short papers presented were thoroughly reviewed and selected from 44 submissions. The papers are organized according to the following topics: ontologies and meta modeling; cyber physical systems and block-chains; recommender systems; machine learning based applications; combinatorial optimization; simulations and deep learning.

Integrates database theory with a practical approach to database design and implementation. From publisher description.

This open access book covers the use of data science, including advanced machine learning, big data analytics, Semantic Web technologies, natural language processing, social media analysis, time series analysis, among others, for applications in economics and finance. In addition, it shows some successful applications of advanced data science solutions used to extract new knowledge from data in order to improve economic forecasting models. The book starts with an introduction on the use of data science technologies in economics and finance and is followed by thirteen chapters showing success stories of the application of specific data science methodologies, touching on particular topics related to novel big data sources and technologies for economic analysis (e.g. social media and news); big data models leveraging on supervised/unsupervised (deep) machine learning; natural language processing to build economic and financial indicators; and forecasting and nowcasting of economic variables through time series analysis. This book is relevant to all stakeholders involved in digital and data-intensive research in economics and finance, helping them to understand the main opportunities and challenges, become familiar with the latest methodological findings, and learn how to use and evaluate the performances of novel tools and frameworks. It primarily targets data scientists and business analysts exploiting data science technologies, and it will also be a useful resource to research students in disciplines and courses related to these topics. Overall, readers will learn modern and effective data science solutions to create tangible innovations for economic and financial applications.

This authoritative text/reference describes the state of the art of fog computing, presenting insights from an international selection of renowned experts. A particular focus is provided on development approaches, architectural mechanisms, and measurement metrics for building smart adaptable environments. The coverage also includes important related topics such as device connectivity, security and interoperability, and communication methods. Topics and features: introduces the core concepts and principles of fog computing, and reviews the latest research and best practice relating to fog/edge environments; discusses the vision for an Internet of Things (IoT) in terms of fog computing and other related distributed computing paradigms, such as cloud computing; presents a survey of the key issues and broader aspects of the fog paradigm, as well as the factors that affect adoption of fog computing; examines frameworks and methodologies for fog-based architecture design, improving performance, and measuring quality of experience; proposes tools and methodologies for analyzing large amounts of sensor data from smart city initiatives; describes approaches for designing robust services, management of data-intensive applications, context-aware data analysis, and vehicular networking; identifies potential future research directions and technological innovations in relation to distributed computing environments such as the IoT. This enlightening volume offers essential perspectives for researchers of distributed computing and computer networking, as well as for advanced undergraduate and graduate students pursuing interests in this area. Professional engineers seeking to enhance security and connectivity in their IoT systems will also find this work to be a valuable reference.

From the Foreword: "Big Data Management and Processing is [a] state-of-the-art book that deals with a wide range of topical themes in the field of

Big Data. The book, which probes many issues related to this exciting and rapidly growing field, covers processing, management, analytics, and applications... [It] is a very valuable addition to the literature. It will serve as a source of up-to-date research in this continuously developing area. The book also provides an opportunity for researchers to explore the use of advanced computing technologies and their impact on enhancing our capabilities to conduct more sophisticated studies." ---Sartaj Sahni, University of Florida, USA "Big Data Management and Processing covers the latest Big Data research results in processing, analytics, management and applications. Both fundamental insights and representative applications are provided. This book is a timely and valuable resource for students, researchers and seasoned practitioners in Big Data fields. --Hai Jin, Huazhong University of Science and Technology, China Big Data Management and Processing explores a range of big data related issues and their impact on the design of new computing systems. The twenty-one chapters were carefully selected and feature contributions from several outstanding researchers. The book endeavors to strike a balance between theoretical and practical coverage of innovative problem solving techniques for a range of platforms. It serves as a repository of paradigms, technologies, and applications that target different facets of big data computing systems. The first part of the book explores energy and resource management issues, as well as legal compliance and quality management for Big Data. It covers In-Memory computing and In-Memory data grids, as well as co-scheduling for high performance computing applications. The second part of the book includes comprehensive coverage of Hadoop and Spark, along with security, privacy, and trust challenges and solutions. The latter part of the book covers mining and clustering in Big Data, and includes applications in genomics, hospital big data processing, and vehicular cloud computing. The book also analyzes funding for Big Data projects.

The Easy, Common-Sense Guide to Solving Real Problems with NoSQL The Mere Mortals ® tutorials have earned worldwide praise as the clearest, simplest way to master essential database technologies. Now, there's one for today's exciting new NoSQL databases. NoSQL for Mere Mortals guides you through solving real problems with NoSQL and achieving unprecedented scalability, cost efficiency, flexibility, and availability. Drawing on 20+ years of cutting-edge database experience, Dan Sullivan explains the advantages, use cases, and terminology associated with all four main categories of NoSQL databases: key-value, document, column family, and graph databases. For each, he introduces pragmatic best practices for building high-value applications. Through step-by-step examples, you'll discover how to choose the right database for each task, and use it the right way. Coverage includes --Getting started: What NoSQL databases are, how they differ from relational databases, when to use them, and when not to Data management principles and design criteria: Essential knowledge for creating any database solution, NoSQL or relational --Key-value databases: Gaining more utility from data structures --Document databases: Schemaless databases, normalization and denormalization, mutable documents, indexing, and design patterns --Column family databases: Google's BigTable design, table design, indexing, partitioning, and Big Data Graph databases: Graph/network modeling, design tips, query methods, and traps to avoid Whether you're a database developer, data modeler, database user, or student, learning NoSQL can open up immense new opportunities. As thousands of database professionals already know, For Mere Mortals is the fastest, easiest route to mastery.

This volume contains the technical papers presented in the workshops, which took place at the 6th European Conference on Service-Oriented and Cloud Computing, ESOC 2017, held in Oslo, Norway, September 2017: First International Workshop on Business Process Management in the Cloud, BPM@Cloud 2017; Third International Workshop on Cloud Adoption and Migration, CloudWays 2017. The 9 full papers were carefully reviewed and selected from 12 submissions. In addition, the volume also contains 8 EU Projects papers, describing projects presented at the European Projects Forum, which took place at ESOC 2017. The papers focus on specific topics in service-oriented and cloud computing domains such as limits and/or advantages of existing cloud solutions, future internet technologies, efficient and adaptive deployment and management of service-based applications across multiple clouds, novel cloud service migration practices and solutions, digitization of enterprises in the cloud computing era, federated cloud networking services.

This two volume set LNCS 9827 and LNCS 9828 constitutes the refereed proceedings of the 27th International Conference on Database and Expert Systems Applications, DEXA 2016, held in Porto, Portugal, September 2016. The 39 revised full papers presented together with 29 short papers were carefully reviewed and selected from 137 submissions. The papers discuss a range of topics including: Temporal, Spatial, and High Dimensional Databases; Data Mining; Authenticity, Privacy, Security, and Trust; Data Clustering; Distributed and Big Data Processing; Decision Support Systems, and Learning; Data Streams; Data Integration, and Interoperability; Semantic Web, and Data Semantics; Social Networks, and Network Analysis; Linked Data; Data Analysis; NoSQL, NewSQL; Multimedia Data; Personal Information Management; Semantic Web and Ontologies; Database and Information System Architectures; Query Answering and Optimization; Information Retrieval, and Keyword Search; Data Modelling, and Uncertainty.

Recent times are witnessing rapid development in machine learning algorithm systems, especially in reinforcement learning, natural language processing, computer and robot vision, image processing, speech, and emotional processing and understanding. In tune with the increasing importance and relevance of machine learning models, algorithms, and their applications, and with the emergence of more innovative uses-cases of deep learning and artificial intelligence, the current volume presents a few innovative research works and their applications in real-world, such as stock trading, medical and healthcare systems, and software automation. The chapters in the book illustrate how machine learning and deep learning algorithms and models are designed, optimized, and deployed. The volume will be useful for advanced graduate and doctoral students, researchers, faculty members of universities, practicing data scientists and data engineers, professionals, and consultants working on the broad areas of machine learning, deep learning, and artificial intelligence.

A hands-on guide to leveraging NoSQL databases NoSQL databases are an efficient and powerful tool for storing and manipulating vast quantities of data. Most NoSQL databases scale well as data grows. In addition, they are often malleable and flexible enough to accommodate semi-structured and sparse data sets. This comprehensive hands-on guide presents fundamental concepts and practical solutions for getting you ready to use NoSQL databases. Expert author Shashank Tiwari begins with a helpful introduction on the subject of NoSQL, explains its characteristics and typical uses, and looks at where it fits in the application stack. Unique insights help you choose which NoSQL solutions are best for solving your specific data storage needs. Professional NoSQL: Demystifies the concepts that relate to NoSQL databases, including column-family oriented stores, key/value databases, and document databases. Delves into installing and configuring a number of NoSQL products and the Hadoop family of products. Explains ways of storing, accessing, and querying data in NoSQL databases through examples that use MongoDB, HBase, Cassandra, Redis, CouchDB, Google App Engine Datastore and more. Looks at architecture and internals. Provides guidelines for optimal usage, performance tuning, and scalable configurations. Presents a number of tools and utilities relating to NoSQL, distributed platforms, and scalable processing, including Hive, Pig, RRDtool, Nagios, and more.

In the last 25 years, information systems have had a disruptive effect on society and business. Up until recently though, the majority of passengers and goods were transported by sea in many ways similar to the way they were at the turn of the previous century. Gradually, advanced information technologies are being introduced, in an attempt to make shipping safer, greener, more efficient, and transparent. The emerging field of Maritime Informatics studies the application of information technology and information systems to maritime transportation. Maritime Informatics can be considered as both a field of study and domain of application. As an application domain, it is the outlet of innovations originating from data science and artificial intelligence; as a field of study, it is positioned between computer science and marine engineering. This new field's complexity lies within this duality because it is faced with disciplinary barriers yet demands a systemic, transdisciplinary approach. At present, there is a growing body of knowledge that remains undocumented in a single source or textbook designed to assist students and practitioners. This highly useful textbook/reference starts by introducing required knowledge, algorithmic approaches, and technical details, before presenting real-world applications. The aim is to present interested audiences with an overview of the main technological innovations having a disruptive effect on the maritime industry, as well as to discuss principal ideas, methods of operation and applications, and future developments. The material in this unique volume provides requisite core knowledge for undergraduate or postgraduate students, employing an analytical approach with numerous real-world examples and case studies.