

Database Principles 2nd Edition

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

Information Modeling and Relational Databases, Second Edition, provides an introduction to ORM (Object-Role Modeling) and much more. In fact, it is the only book to go beyond introductory coverage and provide all of the in-depth instruction you need to transform knowledge from domain experts into a sound database design. This book is intended for anyone with a stake in the accuracy and efficacy of databases: systems analysts, information modelers, database designers and administrators, and programmers. Terry Halpin, a pioneer in the development of ORM, blends conceptual information with practical instruction that will let you begin using ORM effectively as soon as possible. Supported by examples, exercises, and useful background information, his step-by-step approach teaches you to develop a natural-language-based ORM model, and then, where needed, abstract ER and UML models from it. This book will quickly make you proficient in the modeling technique that is proving vital to the development of accurate and efficient databases that best meet real business objectives. Presents the most in-depth coverage of Object-Role Modeling available anywhere, including a thorough update of the book for ORM2, as well as UML2 and E-R (Entity-Relationship) modeling. Includes clear coverage of relational database concepts, and the latest developments in SQL and XML, including a new chapter on the impact of XML on information modeling, exchange and transformation. New and improved case studies and exercises are provided for many topics.

Traditional methods for handling spatial data are encumbered by the assumption of separate origins for horizontal and vertical measurements, but modern measurement systems operate in a 3-D spatial environment. The 3-D Global Spatial Data Model: Principles and Applications, Second Edition maintains a new model for handling digital spatial data, the global spatial data model or GSDM. The GSDM preserves the integrity of three-dimensional spatial data while also providing additional benefits such as simpler equations, worldwide standardization, and the ability to track spatial data accuracy with greater specificity and convenience. This second edition expands to new topics that satisfy a growing need in the GIS, professional surveyor, machine control, and Big Data communities while continuing to embrace the earth center fixed coordinate system as the fundamental point of origin of one, two, and three-dimensional data sets. Ideal for both beginner and advanced levels, this book also provides guidance and insight on how to link to the data collected and stored in legacy systems. The second edition of a comprehensive state-of-the-art graduate level text on microeconomic methods, substantially revised and updated. The second edition of this acclaimed graduate text provides a unified treatment of two methods used in contemporary econometric research, cross section and data panel methods. By focusing on assumptions that can be given behavioral content, the book maintains an appropriate level of rigor while emphasizing intuitive thinking. The analysis covers both linear and nonlinear models, including models with dynamics and/or individual heterogeneity. In addition to general estimation frameworks (particular methods of moments and maximum likelihood), specific linear and nonlinear methods are covered in detail, including probit and logit models and their multivariate, Tobit models, models for count data, censored and missing data schemes, causal (or treatment) effects, and duration analysis. Econometric Analysis of Cross Section and Panel Data was the first graduate econometrics text to focus on microeconomic data structures, allowing assumptions to be separated into population and sampling assumptions. This second edition has been substantially updated and revised. Improvements include a broader class of models for missing data problems; more detailed treatment of cluster problems, an important topic for empirical researchers; expanded discussion of "generalized instrumental variables" (GIV) estimation; new coverage (based on the author's own recent research) of inverse probability weighting; a more complete framework for estimating treatment effects with panel data, and a firmly established link between econometric approaches to nonlinear panel data and the "generalized estimating equation" literature popular in statistics and other fields. New attention is given to explaining when particular econometric methods can be applied; the goal is not only to tell readers what does work, but why certain "obvious" procedures do not. The numerous included exercises, both theoretical and computer-based, allow the reader to extend methods covered in the text and discover new insights.

Beginning Database Design, Second Edition provides short, easy-to-read explanations of how to get database design right the first time. This book offers numerous examples to help you avoid the many pitfalls that entrap new and not-so-new database designers. Through the help of use cases and class diagrams modeled in the UML, you'll learn to discover and represent the details and scope of any design problem you choose to attack. Database design is not an exact science. Many are surprised to find that problems with their databases are caused by poor design rather than by difficulties in using the database management software. Beginning Database Design, Second Edition helps you ask and answer important questions about your data so you can understand the problem you are trying to solve and create a pragmatic design capturing the essentials while leaving the door open for refinements and extension at a later stage. Solid database design principles and examples help demonstrate the consequences of simplifications and pragmatic decisions. The rationale is to try to keep a design simple, but allow room for development as situations change or resources permit. Provides solid design principles by which to avoid pitfalls and support changing needs. Includes numerous examples of good and bad design decisions and their consequences. Shows a modern method for documenting design using the Unified Modeling Language.

Over the past 5 years, the concept of big data has matured, data science has grown exponentially, and data architecture has become a standard part of organizational decision-making. Throughout all this change, the basic principles that shape the architecture of data have remained the same. There remains a need for people to take a look at the "bigger picture" and to understand where their data fit into the grand scheme of things. Data Architecture: A Primer for the Data Scientist, Second Edition addresses the larger architectural picture of how big data fits within the existing information infrastructure or data warehousing systems. This is an essential topic not only for data scientists, analysts, and managers but also for researchers and engineers who increasingly need to deal with large and complex sets of data. Until data are gathered and can be placed into an existing framework or architecture, they cannot be used to their full potential. Drawing upon years of practical experience and using numerous examples and case studies from across various industries, the authors seek to explain this larger picture into which big data fits, giving data scientists the necessary context for how pieces of the puzzle should fit together. New case studies include expanded coverage of textual management and analytics. New chapters on visualization and big data. Discussion of new visualizations of the end-state architecture.

An updated, introductory management book which discusses object oriented data modeling and client server platforms. KEY FEATURES: It explores management and design within the context of the database development life cycle.

A preliminary edition of this book was published from O'Reilly (ISBN 9780596550066). SQLite is a small, embeddable, SQL-based, relational database management system. It has been widely used in low- to medium-tier database applications, especially in embedded devices. This book provides a comprehensive description of SQLite database system. It describes design principles, engineering trade-offs, implementation issues, and operations of SQLite.

During the past decade there has been an explosion in computation and information technology. With it have come vast amounts of data in a variety of fields such as medicine, biology, finance, and marketing. The challenge of understanding these data has led to the development of new tools in the field of statistics, and spawned new areas such as data mining, machine learning, and bioinformatics. Many of these tools have common underpinnings but are often expressed with different terminology. This book describes the important ideas in these areas in a common conceptual framework. While the approach is statistical, the emphasis is

on concepts rather than mathematics. Many examples are given, with a liberal use of color graphics. It should be a valuable resource for statisticians and anyone interested in data mining in science or industry. The book's coverage is broad, from supervised learning (prediction) to unsupervised learning. The many topics include neural networks, support vector machines, classification trees and boosting---the first comprehensive treatment of this topic in any book. This major new edition features many topics not covered in the original, including graphical models, random forests, ensemble methods, least angle regression & path algorithms for the lasso, non-negative matrix factorization, and spectral clustering. There is also a chapter on methods for "wide" data (p bigger than n), including multiple testing and false discovery rates. Trevor Hastie, Robert Tibshirani, and Jerome Friedman are professors of statistics at Stanford University. They are prominent researchers in this area: Hastie and Tibshirani developed generalized additive models and wrote a popular book of that title. Hastie co-developed much of the statistical modeling software and environment in R/S-PLUS and invented principal curves and surfaces. Tibshirani proposed the lasso and is co-author of the very successful *An Introduction to the Bootstrap*. Friedman is the co-inventor of many data-mining tools including CART, MARS, projection pursuit and gradient boosting.

Database Systems: A Pragmatic Approach is a classroom textbook for use by students who are learning about relational databases, and the professors who teach them. It discusses the database as an essential component of a software system, as well as a valuable, mission critical corporate resource. The book is based on lecture notes that have been tested and proven over several years, with outstanding results. It also exemplifies mastery of the technique of combining and balancing theory with practice, to give students their best chance at success. Upholding his aim for brevity, comprehensive coverage, and relevance, author Elvis C. Foster's practical and methodical discussion style gets straight to the salient issues, and avoids unnecessary fluff as well as an overkill of theoretical calculations. The book discusses concepts, principles, design, implementation, and management issues of databases. Each chapter is organized systematically into brief, reader-friendly sections, with itemization of the important points to be remembered. It adopts a methodical and pragmatic approach to solving database systems problems. Diagrams and illustrations also sum up the salient points to enhance learning. Additionally, the book includes a number of Foster's original methodologies that add clarity and creativity to the database modeling and design experience while making a novel contribution to the discipline. Everything combines to make *Database Systems: A Pragmatic Approach* an excellent textbook for students, and an excellent resource on theory for the practitioner.

This invaluable addition to any data scientist's library shows you how to apply the R programming language and useful statistical techniques to everyday business situations as well as how to effectively present results to audiences of all levels. To answer the ever-increasing demand for machine learning and analysis, this new edition boasts additional R tools, modeling techniques, and more. *Practical Data Science with R, Second Edition* takes a practice-oriented approach to explaining basic principles in the ever-expanding field of data science. You'll jump right to real-world use cases as you apply the R programming language and statistical analysis techniques to carefully explained examples based in marketing, business intelligence, and decision support. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. *Database Systems: The Complete Book* is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the latest database standards SQL:1999, SQL/PSM, SQL/CLI, JDBC, ODL, and XML, with broader coverage of SQL than most other texts. The second half of the book provides in-depth coverage of databases from the point of view of the DBMS implementor. It focuses on storage structures, query processing, and transaction management. The book covers the main techniques in these areas with broader coverage of query optimization than most other texts, along with advanced topics including multidimensional and bitmap indexes, distributed transactions, and information integration techniques.

Joe Celko has looked deep into the code of SQL programmers and found a consistent and troubling pattern - a frightening lack of consistency between their individual encoding schemes and those of the industries in which they operate. This translates into a series of incompatible databases, each one an island unto itself that is unable to share information with others in an age of internationalization and business interdependence. Such incompatibility severely hinders information flow and the quality of company data. *Data, Measurements and Standards in SQL* reveals the shift these programmers need to make to overcome this deadlock. By collecting and detailing the diverse standards of myriad industries, and then giving a declaration for the units that can be used in an SQL schema, Celko enables readers to write and implement portable data that can interface to any number of external application systems! This book doesn't limit itself to one subject, but serves as a detailed synopsis of measurement scales and data standards for all industries, thereby giving RDBMS programmers and designers the knowledge and know-how they need to communicate effectively across business boundaries. * Collects and details the diverse data standards of myriad industries under one cover, thereby creating a definitive, one-stop-shopping opportunity for database programmers. * Enables readers to write and implement portable data that can interface to any number external application systems, allowing readers to cross business boundaries and move up the career ladder. * Expert advice from one of the most-read SQL authors in the world who is well known for his ten years of service on the ANSI SQL standards committee and Readers Choice Award winning column in *Intelligent Enterprise*.

The second edition of this innovative work again provides a unique perspective on the clinical discovery process by providing input from experts within the NIH on the principles and practice of clinical research. Molecular medicine, genomics, and proteomics have opened vast opportunities for translation of basic science observations to the bedside through clinical research. As an introductory reference it gives clinical investigators in all fields an awareness of the tools required to ensure research protocols are well designed and comply with the rigorous regulatory requirements necessary to maximize the safety of research subjects. Complete with sections on the history of clinical research and ethics, copious figures and charts, and sample documents it serves as an excellent companion text for any course on clinical research and as a must-have reference for seasoned researchers.

*Incorporates new chapters on Managing Conflicts of Interest in Human Subjects Research, Clinical Research from the Patient's Perspective, The Clinical Researcher and the Media, Data Management in Clinical Research, Evaluation of a Protocol Budget,

Clinical Research from the Industry Perspective, and Genetics in Clinical Research *Addresses the vast opportunities for translation of basic science observations to the bedside through clinical research *Delves into data management and addresses how to collect data and use it for discovery *Contains valuable, up-to-date information on how to obtain funding from the federal government

Database Modeling and Design, Fifth Edition, focuses on techniques for database design in relational database systems. This extensively revised fifth edition features clear explanations, lots of terrific examples and an illustrative case, and practical advice, with design rules that are applicable to any SQL-based system. The common examples are based on real-life experiences and have been thoroughly class-tested. This book is immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data. It is ideal for a stand-alone data management course focused on logical database design, or a supplement to an introductory text for introductory database management. In-depth detail and plenty of real-world, practical examples throughout Loaded with design rules and illustrative case studies that are applicable to any SQL, UML, or XML-based system Immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data.

The book describes the major components of Oracle such as SQL*Plus, PL/SQL, indexing, security and integrity, and distributed databases. Underlying principles are also described: there are chapters on the objectives of database systems and on the relational model. Its broad coverage also includes database design techniques such as normalisation and entity-relationship modelling. The Oracle application development tools: SQL*Forms version 3 and SQL*Reportwriter are covered, since these tools are still widely used in universities and in industry. In preparation for Oracle version 8, object oriented concepts and the object-relational model are described.

Oracle E-Business Suite Controls: Foundational Principles is the second edition book as an expansion of Jeffrey Hare's book titled "Oracle E-Business Suite Controls: Application Security Best Practices" This book provides an overview of the key elements in E-Business Suite that impact security and internal controls with an expanded discussion of best practices.

For programmers who prefer content to frills, this guide has succinct and straightforward information for putting Access to its full, individually tailored use.

Database Principles Fundamentals of Design, Implementation, and Management South Western Educational Publishing
Part I: Process design -- Introduction to design -- Process flowsheet development -- Utilities and energy efficient design -- Process simulation -- Instrumentation and process control -- Materials of construction -- Capital cost estimating -- Estimating revenues and production costs -- Economic evaluation of projects -- Safety and loss prevention -- General site considerations -- Optimization in design -- Part II: Plant design -- Equipment selection, specification and design -- Design of pressure vessels -- Design of reactors and mixers -- Separation of fluids -- Separation columns (distillation, absorption and extraction) -- Specification and design of solids-handling equipment -- Heat transfer equipment -- Transport and storage of fluids.

Forecasting is required in many situations. Stocking an inventory may require forecasts of demand months in advance. Telecommunication routing requires traffic forecasts a few minutes ahead. Whatever the circumstances or time horizons involved, forecasting is an important aid in effective and efficient planning. This textbook provides a comprehensive introduction to forecasting methods and presents enough information about each method for readers to use them sensibly.

*The most updated PostgreSQL book on the market, covering version 8.0 *Highlights the most popular PostgreSQL APIs, including C, Perl, PHP, and Java *This is two books in one; it simultaneously covers key relational database design principles, while teaching PostgreSQL

Principles of Transaction Processing is a comprehensive guide to developing applications, designing systems, and evaluating engineering products. The book provides detailed discussions of the internal workings of transaction processing systems, and it discusses how these systems work and how best to utilize them. It covers the architecture of Web Application Servers and transactional communication paradigms. The book is divided into 11 chapters, which cover the following: Overview of transaction processing application and system structure Software abstractions found in transaction processing systems Architecture of multitier applications and the functions of transactional middleware and database servers Queued transaction processing and its internals, with IBM's Websphere MQ and Oracle's Stream AQ as examples Business process management and its mechanisms Description of the two-phase locking function, B-tree locking and multigranularity locking used in SQL database systems and nested transaction locking System recovery and its failures Two-phase commit protocol Comparison between the tradeoffs of replicating servers versus replication resources Transactional middleware products and standards Future trends, such as cloud computing platforms, composing scalable systems using distributed computing components, the use of flash storage to replace disks and data streams from sensor devices as a source of transaction requests. The text meets the needs of systems professionals, such as IT application programmers who construct TP applications, application analysts, and product developers. The book will also be invaluable to students and novices in application programming. Complete revision of the classic "non mathematical" transaction processing reference for systems professionals. Updated to focus on the needs of transaction processing via the Internet-- the main focus of business data processing investments, via web application servers, SOA, and important new TP standards. Retains the practical, non-mathematical, but thorough conceptual basis of the first edition.

An authoritative guide to the theory and techniques of biomedical informatics.

This book provides a systematic and comparative description of the vast number of research issues related to the quality of data and information. It does so by delivering a sound, integrated and comprehensive overview of the state of the art and future development of data and information quality in databases and information systems. To this end, it presents an extensive

description of the techniques that constitute the core of data and information quality research, including record linkage (also called object identification), data integration, error localization and correction, and examines the related techniques in a comprehensive and original methodological framework. Quality dimension definitions and adopted models are also analyzed in detail, and differences between the proposed solutions are highlighted and discussed. Furthermore, while systematically describing data and information quality as an autonomous research area, paradigms and influences deriving from other areas, such as probability theory, statistical data analysis, data mining, knowledge representation, and machine learning are also included. Last not least, the book also highlights very practical solutions, such as methodologies, benchmarks for the most effective techniques, case studies, and examples. The book has been written primarily for researchers in the fields of databases and information management or in natural sciences who are interested in investigating properties of data and information that have an impact on the quality of experiments, processes and on real life. The material presented is also sufficiently self-contained for masters or PhD-level courses, and it covers all the fundamentals and topics without the need for other textbooks. Data and information system administrators and practitioners, who deal with systems exposed to data-quality issues and as a result need a systematization of the field and practical methods in the area, will also benefit from the combination of concrete practical approaches with sound theoretical formalisms.

Practical and easy to understand Database Principles: Fundamentals of Design, Implementation, and Management, 10/e, International Edition gives readers a solid foundation in database design and implementation. Filled with visual aids such as diagrams, illustrations, and tables, this market-leading book provides in-depth coverage of database design, demonstrating that the key to successful database implementation is in proper design of databases to fit within a larger strategic view of the data environment. Renowned for its clear, straightforward writing style, the tenth edition has been thoroughly updated to include hot topics such as green computing/sustainability for modern data centers, the role of redundant relationships, and examples of web-database connectivity and code security. In addition, new review questions, problem sets, and cases have been added throughout the book so that readers have multiple opportunities to test their understanding and develop real and useful design skills.

The nature of an information system; Naming; Relationships; Attributes; Types and categories and sets; Models; The record model; The other three popular models; The modelling of relationships; Elementary concepts; Philosophy.

This third edition of a classic textbook can be used to teach at the senior undergraduate and graduate levels. The material concentrates on fundamental theories as well as techniques and algorithms. The advent of the Internet and the World Wide Web, and, more recently, the emergence of cloud computing and streaming data applications, has forced a renewal of interest in distributed and parallel data management, while, at the same time, requiring a rethinking of some of the traditional techniques. This book covers the breadth and depth of this re-emerging field. The coverage consists of two parts. The first part discusses the fundamental principles of distributed data management and includes distribution design, data integration, distributed query processing and optimization, distributed transaction management, and replication. The second part focuses on more advanced topics and includes discussion of parallel database systems, distributed object management, peer-to-peer data management, web data management, data stream systems, and cloud computing. New in this Edition: • New chapters, covering database replication, database integration, multidatabase query processing, peer-to-peer data management, and web data management. • Coverage of emerging topics such as data streams and cloud computing • Extensive revisions and updates based on years of class testing and feedback Ancillary teaching materials are available.

Combines language tutorials with application design advice to cover the PHP server-side scripting language and the MySQL database engine.

From a review of the first edition: "Modern Data Science with R... is rich with examples and is guided by a strong narrative voice. What's more, it presents an organizing framework that makes a convincing argument that data science is a course distinct from applied statistics" (The American Statistician). Modern Data Science with R is a comprehensive data science textbook for undergraduates that incorporates statistical and computational thinking to solve real-world data problems. Rather than focus exclusively on case studies or programming syntax, this book illustrates how statistical programming in the state-of-the-art R/RStudio computing environment can be leveraged to extract meaningful information from a variety of data in the service of addressing compelling questions. The second edition is updated to reflect the growing influence of the tidyverse set of packages. All code in the book has been revised and styled to be more readable and easier to understand. New functionality from packages like sf, purrr, tidymodels, and tidytext is now integrated into the text. All chapters have been revised, and several have been split, re-organized, or re-imagined to meet the shifting landscape of best practice.

This book provides a comprehensive and up to date treatment of theory and practical implementation in Register-based statistics. It begins by defining the area, before explaining how to structure such systems, as well as detailing alternative approaches. It explains how to create statistical registers, how to implement quality assurance, and the use of IT systems for register-based statistics. Further to this, clear details are given about the practicalities of implementing such statistical methods, such as protection of privacy and the coordination and coherence of such an undertaking. This edition offers a full understanding of both the principles and practices of this increasingly popular area of statistics, and can be considered a first step to a more systematic way of working with register-statistical issues. This book addresses the growing global interest in the topic and employs a much broader, more international approach than the 1st edition. New chapters explore different kinds of register-based surveys, such as preconditions for register-based statistics and comparing sample survey and administrative data. Furthermore, the authors present discussions on register-based census, national accounts and the transition towards a register-based system as well as presenting new chapters on quality assessment of administrative sources and production process quality.

This book is a comprehensive, practical, and student-friendly textbook addressing fundamental concepts in database design and applications.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. The Language of SQL, Second Edition Many SQL texts attempt to serve as an encyclopedic reference on SQL syntax -- an approach that is often counterproductive, because that information is readily available in online references published by the major database vendors. For SQL beginners, it's more important for a book to focus on general concepts and to offer clear explanations and examples of what various SQL statements can accomplish. This is that book. A number of features make The Language of SQL unique among introductory SQL books. First, you will not be required to download software or sit with a computer as you read the text. The intent of this book is to provide examples of SQL usage that can be

understood simply by reading. Second, topics are organized in an intuitive and logical sequence. SQL keywords are introduced one at a time, allowing you to grow your understanding as you encounter new terms and concepts. Finally, this book covers the syntax of three widely used databases: Microsoft SQL Server, MySQL, and Oracle. Special "Database Differences" sidebars clearly show you any differences in syntax among these three databases, and instructions are included on how to obtain and install free versions of the databases. This is the only book you need to gain a quick working knowledge of SQL and relational databases.

- Learn How To... Use SQL to retrieve data from relational databases
- Apply functions and calculations to data
- Group and summarize data in a variety of useful ways
- Use complex logic to retrieve only the data you need
- Update data and create new tables
- Design relational databases so that data retrieval is easy and intuitive
- Use spreadsheets to transform your data into meaningful displays
- Retrieve data from multiple tables via joins, subqueries, views, and set logic
- Create, modify, and execute stored procedures
- Install Microsoft SQL Server, MySQL, or Oracle

"This book takes the somewhat daunting process of database design and breaks it into completely manageable and understandable components. Mike's approach whilst simple is completely professional, and I can recommend this book to any novice database designer." --Sandra Barker, Lecturer, University of South Australia, Australia

"Databases are a critical infrastructure technology for information systems and today's business. Mike Hernandez has written a literate explanation of database technology--a topic that is intricate and often obscure. If you design databases yourself, this book will educate you about pitfalls and show you what to do. If you purchase products that use a database, the book explains the technology so that you can understand what the vendor is doing and assess their products better." --Michael Blaha, consultant and trainer, author of *A Manager's Guide to Database Technology*

"If you told me that Mike Hernandez could improve on the first edition of *Database Design for Mere Mortals* I wouldn't have believed you, but he did! The second edition is packed with more real-world examples, detailed explanations, and even includes database-design tools on the CD-ROM! This is a must-read for anyone who is even remotely interested in relational database design, from the individual who is called upon occasionally to create a useful tool at work, to the seasoned professional who wants to brush up on the fundamentals. Simply put, if you want to do it right, read this book!" --Matt Greer, Process Control Development, The Dow Chemical Company

"Mike's approach to database design is totally common-sense based, yet he's adhered to all the rules of good relational database design. I use Mike's books in my starter database-design class, and I recommend his books to anyone who's interested in learning how to design databases or how to write SQL queries." --Michelle Poollet, President, MVDS, Inc.

"Slapping together sophisticated applications with poorly designed data will hurt you just as much now as when Mike wrote his first edition, perhaps even more. Whether you're just getting started developing with data or are a seasoned pro; whether you've read Mike's previous book or this is your first; whether you're happier letting someone else design your data or you love doing it yourself--this is the book for you. Mike's ability to explain these concepts in a way that's not only clear, but fun, continues to amaze me." --From the Foreword by Ken Getz, MCW Technologies, coauthor *ASP.NET Developer's JumpStart*

"The first edition of Mike Hernandez's book *Database Design for Mere Mortals* was one of the few books that survived the cut when I moved my office to smaller quarters. The second edition expands and improves on the original in so many ways. It is not only a good, clear read, but contains a remarkable quantity of clear, concise thinking on a very complex subject. It's a must for anyone interested in the subject of database design." --Malcolm C. Rubel, Performance Dynamics Associates

"Mike's excellent guide to relational database design deserves a second edition. His book is an essential tool for fledgling Microsoft Access and other desktop database developers, as well as for client/server pros. I recommend it highly to all my readers." --Roger Jennings, author of *Special Edition Using Access 2002*

"There are no silver bullets! Database technology has advanced dramatically, the newest crop of database servers perform operations faster than anyone could have imagined six years ago, but none of these technological advances will help fix a bad database design, or capture data that you forgot to include! *Database Design for Mere Mortals(TM)*, Second Edition, helps you design your database right in the first place!" --Matt Nunn, Product Manager, SQL Server, Microsoft Corporation

"When my brother started his professional career as a developer, I gave him Mike's book to help him understand database concepts and make real-world application of database technology. When I need a refresher on the finer points of database design, this is the book I pick up. I do not think that there is a better testimony to the value of a book than that it gets used. For this reason I have wholeheartedly recommended to my peers and students that they utilize this book in their day-to-day development tasks." --Chris Kunicki, Senior Consultant, OfficeZealot.com

"Mike has always had an incredible knack for taking the most complex topics, breaking them down, and explaining them so that anyone can 'get it.' He has honed and polished his first very, very good edition and made it even better. If you're just starting out building database applications, this book is a must-read cover to cover. Expert designers will find Mike's approach fresh and enlightening and a source of great material for training others." --John Viescas, President, Viescas Consulting, Inc., author of *Running Microsoft Access 2000* and coauthor of *SQL Queries for Mere Mortals*

"Whether you need to learn about relational database design in general, design a relational database, understand relational database terminology, or learn best practices for implementing a relational database, *Database Design for Mere Mortals(TM)*, Second Edition, is an indispensable book that you'll refer to often. With his many years of real-world experience designing relational databases, Michael shows you how to analyze and improve existing databases, implement keys, define table relationships and business rules, and create data views, resulting in data integrity, uniform access to data, and reduced data-entry errors." --Paul Cornell, Site Editor, MSDN Office Developer Center

Sound database design can save hours of development time and ensure functionality and reliability. *Database Design for Mere Mortals(TM)*, Second Edition, is a straightforward, platform-independent tutorial on the basic principles of relational database design. It provides a commonsense design methodology for developing databases that work. Database design expert Michael J. Hernandez has expanded his best-selling first edition, maintaining its hands-on approach and accessibility while updating its coverage and including even more examples and illustrations. This edition features a CD-ROM that includes diagrams of sample databases, as well as design guidelines, documentation forms, and examples of the database design process. This book will give you the knowledge and tools you need to create efficient and effective relational databases.

Principles of Big Data helps readers avoid the common mistakes that endanger all Big Data projects. By stressing simple, fundamental concepts, this book teaches readers how to organize large volumes of complex data, and how to achieve data permanence when the content of the data is constantly changing. General methods for data verification and validation, as specifically applied to Big Data resources, are stressed throughout the book. The book demonstrates how adept analysts can find relationships among data objects held in disparate Big Data resources, when the data objects are endowed with semantic support (i.e., organized in classes of uniquely identified data objects). Readers will learn how their data can be integrated with data from

other resources, and how the data extracted from Big Data resources can be used for purposes beyond those imagined by the data creators. Learn general methods for specifying Big Data in a way that is understandable to humans and to computers Avoid the pitfalls in Big Data design and analysis Understand how to create and use Big Data safely and responsibly with a set of laws, regulations and ethical standards that apply to the acquisition, distribution and integration of Big Data resources SQL is full of difficulties and traps for the unwary. You can avoid them if you understand relational theory, but only if you know how to put the theory into practice. In this insightful book, author C.J. Date explains relational theory in depth, and demonstrates through numerous examples and exercises how you can apply it directly to your use of SQL. This second edition includes new material on recursive queries, "missing information" without nulls, new update operators, and topics such as aggregate operators, grouping and ungrouping, and view updating. If you have a modest-to-advanced background in SQL, you'll learn how to deal with a host of common SQL dilemmas. Why is proper column naming so important? Nulls in your database are causing you to get wrong answers. Why? What can you do about it? Is it possible to write an SQL query to find employees who have never been in the same department for more than six months at a time? SQL supports "quantified comparisons," but they're better avoided. Why? How do you avoid them? Constraints are crucially important, but most SQL products don't support them properly. What can you do to resolve this situation? Database theory and practice have evolved since the relational model was developed more than 40 years ago. SQL and Relational Theory draws on decades of research to present the most up-to-date treatment of SQL available. C.J. Date has a stature that is unique within the database industry. A prolific writer well known for the bestselling textbook An Introduction to Database Systems (Addison-Wesley), he has an exceptionally clear style when writing about complex principles and theory.

This lean, focused text concentrates on giving students a clear understanding of database fundamentals while providing a broad survey of all the major topics of the field. The result is a text that is easily covered in one semester, and that only includes topics relevant to the database course. Mark Gillenson, an associate editor of the Journal of Database Management, has 15 years experience of working with and teaching at IBM Corp. and 15 years of teaching experience at the college level. He writes in a clear, friendly style that progresses step-by-step through all of the major database topics. Each chapter begins with a story about a real company's database application, and is packed with examples. When students finish the text, they will be able to immediately apply what they've learned in business.

Organisations increasingly view data as a valuable corporate asset and its effective management can be vital to an organisation's success. This professional reference guide covers all the key areas including database development, data quality and corporate data modelling. It is not based on a particular proprietary system; it is business focused, providing the knowledge and techniques required to successfully implement the data management function.

Geared to IT professionals eager to get into the all-important field of data warehousing, this book explores all topics needed by those who design and implement data warehouses. Readers will learn about planning requirements, architecture, infrastructure, data preparation, information delivery, implementation, and maintenance. They'll also find a wealth of industry examples garnered from the author's 25 years of experience in designing and implementing databases and data warehouse applications for major corporations. Market: IT Professionals, Consultants.

[Copyright: 0fced262d46714fafdc8b3b7dbb87d37](#)