Matrix Computations Golub Van Loan 4th Edition

A second course in linear algebra for undergraduates in mathematics, computer science, physics, statistics, and the biological sciences.

Numerical linear algebra, digital signal processing, and parallel algorithms are three disciplines with a great deal of activity in the last few years. The interaction between them has been growing to a level that merits an Advanced Study Institute dedicated to the three areas together. This volume gives an account of the main results in this interdisciplinary field. The following topics emerged as major themes of the meeting: - Singular value and eigenvalue decompositions, including applications, - Toeplitz matrices, including special algorithms and architectures, - Recursive least squares in linear algebra, digital signal processing and control, - Updating and downdating techniques in linear algebra and signal processing, - Stability and sensitivity analysis of special recursive least squares problems, - Special architectures for linear algebra and signal processing. This book contains tutorials on these topics given by leading scientists in each of the three areas. A consider- able number of new research results are presented in contributed papers. The tutorials and papers will be of value to anyone interested in the three disciplines.

Eigenvalue computations are ubiquitous in science and engineering. John Francis's implicitly shifted QR algorithm has been the method of choice for small to medium sized eigenvalue problems since its invention in 1959. This book presents a new view of this classical algorithm. While Francis's original procedure chases bulges, the new version chases core transformations, which allows the development of fast algorithms for eigenvalue problems with a variety of special structures. This also leads to a fast and backward stable algorithm for computing the roots of a polynomial by solving the companion matrix eigenvalue problem. The authors received a SIAM Outstanding Paper prize for this work. This book will be of interest to researchers in numerical linear algebra and their students.

Using an approach that author Alan Laub calls "matrix analysis for grown-ups," this new textbook introduces fundamental concepts of numerical linear algebra and their application to solving certain numerical problems arising in state-space control and systems theory. It is written for advanced undergraduate and beginning graduate students and can be used as a follow-up to Matrix Analysis for Scientists and Engineers (SIAM, 2005), a compact single-semester introduction to matrix analysis for engineers and computational scientists by the same author. Computational Matrix Analysis provides readers with a one-semester introduction to numerical linear algebra; an introduction to statistical condition estimation in book form for the first time; and an overview of certain computational problems in control and systems theory. The book features a number of elements designed to help students learn to use numerical linear algebra in day-to-day computing or research, including a brief review of matrix analysis, including notation, and an introduction to finite (IEEE) arithmetic; discussion and examples of conditioning, stability, and rounding analysis; an introduction to mathematical software topics related to numerical linear algebra; a thorough introduction to Gaussian elimination, along with condition estimation techniques; coverage of linear least squares, with orthogonal reduction and QR factorization; variants of the QR algorithm; and applications of the discussed algorithms.

This is part two of a two-volume book on real analysis and is intended for senior undergraduate students of mathematics who have already been exposed to calculus. The emphasis is on rigour and foundations of analysis. Beginning with the construction of the number systems and set theory, the book discusses the basics of analysis (limits, series, continuity, differentiation, Riemann integration), through to power series, several variable calculus and Fourier analysis, and then finally the Lebesgue integral. These are almost entirely set in the concrete setting of the real line and Euclidean spaces, although there is some material on abstract metric and topological spaces. The book also has appendices on mathematical logic and the decimal system. The entire text (omitting some less central topics) can be taught in two quarters of 25–30 lectures each. The course material is deeply intertwined with the exercises, as it is intended that the student actively learn the material (and practice thinking and writing rigorously) by proving several of the key results in the theory.

This comprehensive textbook is designed for first-year graduate students from a variety of engineering and scientific disciplines.

Linear algebra and matrix theory are fundamental tools in mathematical and physical science, as well as fertile fields for research. This second edition of this acclaimed text presents results of both classic and recent matrix analysis using canonical forms as a unifying theme and demonstrates their importance in a variety of applications. This thoroughly revised and updated second edition is a text for a second course on linear algebra and has more than 1,100 problems and exercises, new sections on the singular value and CS decompositions and the Weyr canonical form, expanded treatments of inverse problems and of block matrices, and much more.

This revised edition provides the mathematical background and algorithmic skills required for the production of numerical software. It includes rewritten and clarified proofs and derivations, as well as new topics such as Arnoldi iteration, and domain decomposition methods.

This is a textbook on classical polynomial and rational approximation theory for the twenty-first century. Aimed at advanced undergraduates and graduate students across all of applied mathematics, it uses MATLAB to teach the field's most important ideas and results. Approximation Theory and Approximation Practice, Extended Edition differs fundamentally from other works on approximation theory in a number of ways: its emphasis is on topics close to numerical algorithms; concepts are illustrated with Chebfun; and each chapter is a PUBLISHable MATLAB M-file, available online. The book centers on theorems and methods for analytic functions, which appear so often in applications, rather than on functions at the edge of discontinuity with their seductive theoretical challenges. Original sources are cited rather than textbooks, and each item in the bibliography is accompanied by an editorial comment. In addition, each chapter has a collection of exercises, which span a wide range from mathematical theory to Chebfun-based numerical experimentation. This textbook is appropriate for advanced undergraduate or graduate students who have an understanding of numerical analysis and complex analysis. It is also appropriate for seasoned mathematicians who use MATLAB.

A concise, insightful, and elegant introduction to the field of numerical linear algebra. Designed for use as a stand-alone textbook in a one-semester, graduate-level course in the topic, it has already been class-tested by MIT and Cornell graduate students from all fields of mathematics, engineering, and the physical sciences. The authors' clear, inviting style and evident love of the field, along with their eloquent presentation of the most fundamental ideas in numerical linear algebra, make it popular with teachers and students alike.

This is the most authoritative and accessible single-volume reference book on applied mathematics. Featuring numerous entries by leading experts and organized thematically, it introduces readers to applied mathematics and its uses; explains key concepts; describes important equations, laws, and functions; looks at exciting areas of research; covers modeling and simulation; explores areas of application; and more. Modeled on the popular Princeton Companion to Mathematics, this volume is an indispensable resource for undergraduate and graduate students, researchers, and practitioners in other disciplines seeking a user-friendly reference book on applied mathematics. Features nearly 200 entries organized thematically and written by an international team of distinguished contributors Presents the major ideas and branches of applied mathematics in a clear and accessible way Explains important mathematical concepts, methods, equations, and applications Introduces the language of applied mathematics and the goals of applied mathematical research Gives a wide range of examples of mathematical modeling Covers continuum mechanics, dynamical systems, numerical analysis, discrete and combinatorial mathematics, mathematical physics, and much more Explores the connections between applied mathematics and other disciplines Includes suggestions for further reading, cross-references, and a comprehensive index

The book's most important contribution is to collect, organize, and explain the many theorems on partially ordered sets in a way that makes them available to the widest possible audience.Chapters: Introduction to Dimension; Crowns, Splits, Stacks, Sums and Products; Characterization Problems for Posets, Lattices, Graphs, and Families of Sets; Hypergraph Coloring, Computational Complexity, and Irreducible Posets; Planar Posets and Trees; Planar Graphs, Planar Maps and Convex Polytopes; Probabilistic Methods in Dimension Theory; Interval and Geometric Containment Orders; Greedy Dimension, Back-Tracking, and Depth First Search; Products of Chains of Bounded Length; Large Minimal Realizers

Describes basic programming principles and their step-by- step applications.Numerous examples are included.

Building on the foundations of its predecessor volume, Matrix Analysis, this book treats in detail several topics in matrix theory not included in the previous volume, but with important applications and of special mathematical interest. As with the previous volume, the authors assume a background knowledge of elementary linear algebra and rudimentary analytical concepts. Many examples and exercises of varying difficulty are included.

Proceedings of the NATO Advanced Study Institute, Leuven, Belgium, August 3-14, 1992

Unique in content and approach, this book covers all the topics that are usually covered in an introduction to scientific computing--but folds in graphics and matrix-vector manipulation in a way that gets readers to appreciate theconnection between continuous mathematics and computing. MATLAB 5 is usedthroughout to encourage experimentation, and each chapter focuses on a different important theorem--allowing readers to appreciate the rigorous side of scientific computing. In addition to standard topical coverage, each chapter includes 1) a sketch of a “hard” problem that involves ill-conditioning, high dimension, etc.; 2)at least one theorem with both a rigorous proof and a “proof by MATLAB” experiment to bolster intuition; 3)at least one recursive algorithm; and 4)at least one connection to a real-world application. The book revolves around examples that are packaged in 200+ M-files, which, collectively, communicate all the key mathematical ideas and an appreciation for the subtleties of numerical computing.Power Tools of the Trade. Polynomial Interpolation. Piecewise Polynomial Interpolation. Numerical Integration. Matrix Computations. Linear Systems. The QR and Cholesky Factorizations. Nonlinear Equations and Optimization. The Initial Value Problem.For engineers and mathematicians.

Revised and updated, the third edition of Golub and Van Loan's classic text in computer science provides essential information about the mathematical background and algorithmic skills required for the production of numerical software. This new edition includes thoroughly revised chapters on matrix multiplication problems and parallel matrix computations, expanded treatment of CS decomposition, an updated overview of floating point arithmetic, a more accurate rendition of the modified Gram-Schmidt process, and new material devoted to GMRES, QMR, and other methods designed to handle the sparse unsymmetric linear system problem.

Accuracy and Stability of Numerical Algorithms gives a thorough, up-to-date treatment of the behavior of numerical algorithms in finite precision arithmetic. It combines algorithmic derivations, perturbation theory, and rounding error analysis, all enlivened by historical perspective and informative quotations. This second edition expands and updates the coverage of the first edition (1996) and includes numerous improvements to the original material. Two new chapters treat symmetric indefinite systems and skew-symmetric systems, and nonlinear systems and Newton's method. Twelve new sections include coverage of additional error bounds for Gaussian elimination, rank revealing LU factorizations, weighted and constrained least squares problems, and the fused multiply-add operation found on some modern computer architectures.

Matrix ComputationsJHU Press

This computationally oriented book describes and explains the mathematical relationships among matrices, moments, orthogonal polynomials, quadrature rules, and the Lanczos and conjugate gradient algorithms. The book bridges different mathematical areas to obtain algorithms to estimate bilinear forms involving two vectors and a function of the matrix. The first part of the book provides the necessary mathematical background and explains the theory. The second part describes the applications and gives numerical examples of the algorithms and techniques developed in the first part. Applications addressed in the book include computing elements of functions of matrices; obtaining estimates of the error norm in iterative methods for solving linear systems and computing parameters in least squares and total least squares; and solving ill-posed problems using Tikhonov regularization. This book will interest researchers in numerical linear algebra and matrix computations, as well as scientists and engineers working on problems involving computation of bilinear forms.

This classic volume covers the fundamentals of two closely related topics: linear systems (linear equations and least-squares) and linear programming (optimizing a linear function subject to linear constraints). For each problem class, stable and efficient numerical algorithms intended for a finite-precision environment are derived and analyzed. While linear algebra and optimization have made huge advances since this book first appeared in 1991, the fundamental principles have not changed. These topics were rarely taught with a unified perspective, and, somewhat surprisingly, this remains true 30 years later. As a result, some of the material in this book can be difficult to find elsewhere—in particular, techniques for updating the LU factorization, descriptions of the simplex method applied to all-inequality form, and the analysis of what happens when using an approximate inverse to solve Ax=b. Numerical Linear Algebra and Optimization is primarily a reference for students who want to learn about numerical techniques for solving linear systems and/or linear programming using the simplex method; however, Chapters 6, 7, and 8 can be used as the text for an upper-division course on linear least squares and linear programming. Understanding is enhanced by numerous exercises.

Matrix functions and matrix equations are widely used in science, engineering and social sciences due to the succinct and insightful way in which they allow problems to be formulated and solutions to be expressed. This book covers materials relevant to advanced undergraduate and graduate courses in numerical linear algebra and scientific computing. It is also well-suited for self-study. The broad content makes it convenient as a general reference to the subjects. Contents:Matrix Functions: A Short Course (Nicholas J Higham and Lijing Lin)A Short Course on Exponential Integrators (Marlis Hochbruck)Matrix Equations and Model Reduction (Peter Benner, Tobias Breiten and Lihong Feng)Rayleigh Quotient Based Optimization Methods for Eigenvalue Problems (Ren-Cang Li)Factorization-Based Sparse Solvers and Preconditioners (Xiaoye Sherry Li) Readership: Researchers and graduate students in numerical and computational mathematics. Key Features:The book covers underlying theory and a variety of algorithms for matrix functions and matrix equations. The book also covers high performance linear system solvers and eigenvalue computations which are computational kernels to matrix functions and matrix equationsThe book provides the current developments and applications beyond the material found in regular university courses and textbooks. It includes a comprehensive list of latest referencesThe authors of the chapters are leading experts who are also well-known for their expository skillsKeywords:Matrix Functions;Matrix Equations;Numerical Linear Algebra;Exponential Integral

Focusing on special matrices and matrices which are in some sense `near' to structured matrices, this volume covers a broad range of topics of current interest in numerical linear algebra. Exploitation of these less obvious structural properties can be of great importance in the design of efficient numerical methods, for example algorithms for matrices with low-rank block structure, matrices with decay, and structured tensor computations. Applications range from quantum chemistry to queuing theory. Structured matrices arise frequently in applications. Examples include banded and sparse matrices, Toeplitz-type matrices, and matrices with semi-separable or quasi-separable structure, as well as Hamiltonian and symplectic matrices. The associated literature is enormous, and many efficient algorithms have been developed for solving problems involving such matrices. The text arose from a C.I.M.E. course held in Cetraro (Italy) in June 2015 which aimed to present this fast growing field to young researchers, exploiting the expertise of five leading lecturers with different theoretical and application perspectives.

The book concludes with a survey of recent developments on coloring graphs on surfaces.

An introduction to computer-based problem-solving using the MATLAB® environment for undergraduates.

A knowledge of matrix algebra is a prerequisite for the study of much of modern statistics, especially the areas of linear statistical models and multivariate statistics. This reference book provides the background in matrix algebra necessary to do research and understand the results in these areas. Essentially self-contained, the book is best-suited for a reader who has had some previous exposure to matrices. Solultions to the exercises are available in the author's "Matrix Algebra: Exercises and Solutions."

This much-needed work presents, among other things, the relevant aspects of the theory of matrix algebra for applications in statistics. Written in an informal style, it addresses computational issues and places more emphasis on applications than existing texts.

Matrix Analysis presents the classical and recent results for matrix analysis that have proved to be important to applied mathematics.

Full of features and applications, this acclaimed textbook for upper undergraduate level and graduate level students includes all the major topics of computational linear algebra, including solution of a system of linear equations, least-squares solutions of linear systems, computation of eigenvalues, eigenvectors, and singular value problems. Drawing from numerous disciplines of science and engineering, the author covers a variety of motivating applications. When a physical problem is posed, the scientific and engineering significance of the solution is clearly stated. Each chapter contains a summary of the important concepts developed in that chapter, suggestions for further reading, and numerous exercises, both theoretical and MATLAB and MATCOM based. The author also provides a list of key words for quick reference. The MATLAB toolkit available online, 'MATCOM', contains implementations of the major algorithms in the book and will enable students to study different algorithms for the same problem, comparing efficiency, stability, and accuracy.

Signal processing applications have burgeoned in the past decade. During the same time, signal processing techniques have matured rapidly and now include tools from many areas of mathematics, computer science, physics, and engineering. This trend will continue as many new signal processing applications are opening up in consumer products and communications systems. In particular, signal processing has been making increasingly sophisticated use of linear algebra on both theoretical and algorithmic fronts. This volume gives particular emphasis to exposing broader contexts of the signal processing problems so that the impact of algorithms and hardware can be better understood; it brings together the writings of signal processing engineers, computer engineers, and applied linear algebraists in an exchange of problems, theories, and techniques. This volume will be of interest to both applied mathematicians and engineers.

Table of contents

This self-contained introduction to numerical linear algebra provides a comprehensive, yet concise, overview of the subject. It includes standard material such as direct methods for solving linear systems and least-squares problems, error, stability and conditioning, basic iterative methods and the calculation of eigenvalues. Later chapters cover more advanced material, such as Krylov subspace methods, multigrid methods, domain decomposition methods, multipole expansions, hierarchical matrices and compressed sensing. The book provides rigorous mathematical proofs throughout, and gives algorithms in general-purpose language-independent form. Requiring only a solid knowledge in linear algebra and basic analysis, this book will be useful for applied mathematicians, engineers, computer scientists, and all those interested in efficiently solving linear problems.

This textbook introduces linear algebra and optimization in the context of machine learning. Examples and exercises are provided throughout this text book together with access to a solution’s manual. This textbook targets graduate level students and professors in computer science, mathematics and data science. Advanced undergraduate students can also use this textbook. The chapters for this textbook are organized as follows: 1. Linear algebra and its applications: The chapters focus on the basics of linear algebra together with their common applications to singular value decomposition, matrix factorization, similarity matrices (kernel methods), and graph analysis. Numerous machine learning applications have been used as examples, such as spectral clustering, kernel-based classification, and outlier detection. The tight integration of linear algebra methods with examples from machine learning differentiates this book from generic volumes on linear algebra. The focus is clearly on the most relevant aspects of linear algebra for machine learning and to teach readers how to apply these concepts. 2. Optimization and its applications: Much of machine learning is posed as an optimization problem in which we try to maximize the accuracy of regression and classification models. The “parent problem” of optimization-centric machine learning is least-squares regression. Interestingly, this problem arises in both linear algebra and optimization, and is one of the key connecting problems of the two fields. Least-squares regression is also the starting point for support vector machines, logistic regression, and recommender systems. Furthermore, the methods for dimensionality reduction and matrix factorization also require the development of optimization methods. A general view of optimization in computational graphs is discussed together with its applications to back propagation in neural networks. A frequent challenge faced by beginners in machine learning is the extensive background required in linear algebra and optimization. One problem is that the existing linear algebra and optimization courses are not specific to machine learning; therefore, one would typically have to complete more course material than is necessary to pick up machine learning. Furthermore, certain types of ideas and tricks from optimization and linear algebra recur more frequently in machine learning than other application-centric settings. Therefore, there is significant value in developing a view of linear algebra and optimization that is better suited to the specific perspective of machine learning.

"An invaluable reference book that should be in every university library." -- Image: Bulletin of the International Linear Algebra Society

An in-depth, theoretical discussion of the two most important classes of algorithms for solving matrix eigenvalue problems.

This is the second volume in a projected five-volume survey of numerical linear algebra and matrix algorithms. It treats the numerical solution of dense and large-scale eigenvalue problems with an emphasis on algorithms and the theoretical background required to understand them. The notes and reference sections contain pointers to other methods along with historical comments. The book is divided into two parts: dense eigenproblems and large eigenproblems. The first part gives a full treatment of the widely used QR algorithm, which is then applied to the solution of generalized eigenproblems and the computation of the singular value decomposition. The second part treats Krylov sequence methods such as the Lanczos and Arnoldi algorithms and presents a new treatment of the Jacobi-Davidson method. These volumes are not intended to be encyclopedic, but provide the reader with the theoretical and practical background to read the research literature and implement or modify new algorithms.

Matrix analysis presented in the context of numerical computation at a basic level.