

# Computational Science And Engineering Gilbert Strang

Nonnegative matrix factorization (NMF) in its modern form has become a standard tool in the analysis of high-dimensional data sets. This book provides a comprehensive and up-to-date account of the most important aspects of the NMF problem and is the first to detail its theoretical aspects, including geometric interpretation, nonnegative rank, complexity, and uniqueness. It explains why understanding these theoretical insights is key to using this computational tool effectively and meaningfully. Nonnegative Matrix Factorization is accessible to a wide audience and is ideal for anyone interested in the workings of NMF. It discusses some new results on the nonnegative rank and the identifiability of NMF and makes available MATLAB codes for readers to run the numerical examples presented in the book. Graduate students starting to work on NMF and researchers interested in better understanding the NMF problem and how they can use it will find this book useful. It can be used in advanced undergraduate and graduate-level courses on numerical linear algebra and on advanced topics in numerical linear algebra and requires only a basic knowledge of linear algebra and optimization.

Computational Methods in Engineering brings to light the numerous uses of numerical methods in engineering. It clearly explains the application of these methods mathematically and practically, emphasizing programming aspects when appropriate. By approaching the cross-disciplinary topic of numerical methods with a flexible approach, Computational Methods in Engineering encourages a well-rounded understanding of the subject. This book's teaching goes beyond the text—detailed exercises (with solutions), real examples of numerical methods in real engineering practices, flowcharts, and MATLAB codes all help you learn the methods directly in the medium that suits you best. Balanced discussion of mathematical principles and engineering applications Detailed step-by-step exercises and practical engineering examples to help engineering students and other readers fully grasp the concepts Concepts are explained through flowcharts and simple MATLAB codes to help you develop additional programming skills

The Portable, Extensible Toolkit for Scientific Computation (PETSc) is an open-source library of advanced data structures and methods for solving linear and nonlinear equations and for managing discretizations. This book uses these modern numerical tools to demonstrate how to solve nonlinear partial differential equations (PDEs) in parallel. It starts from key mathematical concepts, such as Krylov space methods, preconditioning, multigrid, and Newton's method. In PETSc these components are composed at run time into fast solvers. Discretizations are introduced from the beginning, with an emphasis on finite difference and finite element methodologies. The example C programs of the first 12 chapters, listed on the inside front cover, solve (mostly) elliptic and parabolic PDE problems. Discretization leads to large, sparse, and generally nonlinear systems of algebraic equations. For such problems, mathematical solver concepts are explained and illustrated through the examples, with sufficient context to speed further development. PETSc for Partial Differential Equations addresses both discretizations and fast solvers for PDEs, emphasizing practice more than theory. Well-structured examples lead to run-time choices that result in high solver performance and parallel scalability. The last two chapters build on the reader's understanding of fast solver concepts when applying the Firedrake Python finite element solver library. This textbook, the first to cover PETSc programming for nonlinear PDEs, provides an on-ramp for graduate students and researchers to a major area of high-performance computing for science and engineering. It is suitable as a supplement for courses in scientific computing or numerical methods for differential equations.

Multivariable Calculus with Mathematica is a textbook addressing the calculus of several variables. Instead of just using Mathematica to directly solve problems, the students are encouraged to learn the syntax and to write their own code to solve problems. This not only encourages scientific computing skills but at the same time stresses the complete understanding of the mathematics. Questions are provided at the end of the chapters to test the student's theoretical understanding of the mathematics, and there are also computer algebra questions which test the student's ability to apply their knowledge in non-trivial ways. Features Ensures that students are not just using the package to directly solve problems, but learning the syntax to write their own code to solve problems Suitable as a main textbook for a Calculus III course, and as a supplementary text for topics scientific computing, engineering, and mathematical physics Written in a style that engages the students' interest and encourages the understanding of the mathematical ideas

This book describes fundamental computational methods for image reconstruction in computed tomography (CT) with a focus on a pedagogical presentation of these methods and their underlying concepts. Insights into the advantages, limitations, and theoretical and computational aspects of the methods are included, giving a balanced presentation that allows readers to understand and implement CT reconstruction algorithms. Unique in its emphasis on the interplay between modeling, computing, and algorithm development, *Computed Tomography: Algorithms, Insight, and Just Enough Theory* develops the mathematical and computational aspects of three main classes of reconstruction methods: classical filtered back-projection, algebraic iterative methods, and variational methods based on nonlinear numerical optimization algorithms. It spotlights the link between CT and numerical methods, which is rarely discussed in current literature, and describes the effects of incomplete data using both microlocal analysis and singular value decomposition (SVD). This book sets the stage for further exploration of CT algorithms. Readers will be able to grasp the underlying mathematical models to motivate and derive the basic principles of CT reconstruction and will gain basic understanding of fundamental computational challenges of CT, such as the influence of noisy and incomplete data, as well as the reconstruction capabilities and the convergence of the iterative algorithms. Exercises using MATLAB are included, allowing readers to experiment with the algorithms and making the book suitable for teaching and self-study. *Computed Tomography: Algorithms, Insight, and Just Enough Theory* is primarily aimed at students, researchers, and practitioners interested in the computational aspects of X-ray CT and is also relevant for anyone working with other forms of tomography, such as neutron and electron tomography, that share the same mathematical formulation. With its basis in lecture notes developed for a PhD course, it is appropriate as a textbook for courses on computational methods for X-ray CT and computational methods for inverse problems.

Renowned applied mathematician Gilbert Strang teaches applied mathematics with the clear explanations, examples and insights of an experienced teacher. This book progresses steadily through a range of topics from symmetric linear systems to differential equations to least squares and Kalman filtering and optimization. It clearly demonstrates the power of matrix algebra in engineering problem solving. This is an ideal book (beloved by many readers) for a first course on applied mathematics and a reference for more advanced applied mathematicians. The only prerequisite is a basic course in linear algebra.

This book is about computational methods based on operator splitting. It consists of twenty-three chapters written by recognized splitting method contributors and practitioners, and covers a vast spectrum of topics and application areas, including computational mechanics, computational physics, image processing, wireless communication, nonlinear optics, and finance. Therefore, the book presents very versatile aspects of splitting methods and their applications, motivating the cross-fertilization of ideas.

*Applied Mathematics in Engineering and Reliability* contains papers presented at the

International Conference on Applied Mathematics in Engineering and Reliability (ICAMER 2016, Ho Chi Minh City, Viet Nam, 4-6 May 2016). The book covers a wide range of topics within mathematics applied in reliability, risk and engineering, including:- Risk and Relia

Combinatorial Scientific Computing explores the latest research on creating algorithms and software tools to solve key combinatorial problems on large-scale high-performance computing architectures. It includes contributions from international researchers who are pioneers in designing software and applications for high-performance computing systems. The book offers a state-of-the-art overview of the latest research, tool development, and applications. It focuses on load balancing and parallelization on high-performance computers, large-scale optimization, algorithmic differentiation of numerical simulation code, sparse matrix software tools, and combinatorial challenges and applications in large-scale social networks. The authors unify these seemingly disparate areas through a common set of abstractions and algorithms based on combinatorics, graphs, and hypergraphs. Combinatorial algorithms have long played a crucial enabling role in scientific and engineering computations and their importance continues to grow with the demands of new applications and advanced architectures. By addressing current challenges in the field, this volume sets the stage for the accelerated development and deployment of fundamental enabling technologies in high-performance scientific computing. This book differs from traditional numerical analysis texts in that it focuses on the motivation and ideas behind the algorithms presented rather than on detailed analyses of them. It presents a broad overview of methods and software for solving mathematical problems arising in computational modeling and data analysis, including proper problem formulation, selection of effective solution algorithms, and interpretation of results. In the 20 years since its original publication, the modern, fundamental perspective of this book has aged well, and it continues to be used in the classroom. This Classics edition has been updated to include pointers to Python software and the Chebfun package, expansions on barycentric formulation for Lagrange polynomial interpretation and stochastic methods, and the availability of about 100 interactive educational modules that dynamically illustrate the concepts and algorithms in the book. Scientific Computing: An Introductory Survey, Second Edition is intended as both a textbook and a reference for computationally oriented disciplines that need to solve mathematical problems.

Many physical, chemical, biomedical, and technical processes can be described by partial differential equations or dynamical systems. In spite of increasing computational capacities, many problems are of such high complexity that they are solvable only with severe simplifications, and the design of efficient numerical schemes remains a central research challenge. This book presents a tutorial introduction to recent developments in mathematical methods for model reduction and approximation of complex systems. Model Reduction and Approximation: Theory and Algorithms contains three parts that cover (I) sampling-based methods, such as the reduced basis method and proper orthogonal decomposition, (II) approximation of high-dimensional problems by low-rank tensor techniques, and (III) system-theoretic methods, such as balanced truncation, interpolatory methods, and the Loewner framework. It is tutorial in nature, giving an accessible introduction to state-of-the-art model reduction and approximation methods. It also covers a wide range of methods drawn from typically distinct communities (sampling based, tensor based, system-theoretic).?? This book is intended for researchers interested in model reduction and approximation, particularly graduate students and young researchers.

Linear algebra and the foundations of deep learning, together at last! From Professor Gilbert Strang, acclaimed author of Introduction to Linear Algebra, comes Linear Algebra and Learning from Data, the first textbook that teaches linear algebra together with deep learning and neural nets. This readable yet rigorous textbook contains a complete course in the linear algebra and related mathematics that students need to know to get to grips with learning from data.

Included are: the four fundamental subspaces, singular value decompositions, special matrices, large matrix computation techniques, compressed sensing, probability and statistics, optimization, the architecture of neural nets, stochastic gradient descent and backpropagation. Based on a course developed by the author, Introduction to High Performance Scientific Computing introduces methods for adding parallelism to numerical methods for solving differential equations. It contains exercises and programming projects that facilitate learning as well as examples and discussions based on the C programming language, with additional comments for those already familiar with C++. The text provides an overview of concepts and algorithmic techniques for modern scientific computing and is divided into six self-contained parts that can be assembled in any order to create an introductory course using available computer hardware. Part I introduces the C programming language for those not already familiar with programming in a compiled language. Part II describes parallelism on shared memory architectures using OpenMP. Part III details parallelism on computer clusters using MPI for coordinating a computation. Part IV demonstrates the use of graphical programming units (GPUs) to solve problems using the CUDA language for NVIDIA graphics cards. Part V addresses programming on GPUs for non-NVIDIA graphics cards using the OpenCL framework. Finally, Part VI contains a brief discussion of numerical methods and applications, giving the reader an opportunity to test the methods on typical computing problems. The articles that comprise this distinguished annual volume for the Advances in Mechanics and Mathematics series have been written in honor of Gilbert Strang, a world renowned mathematician and exceptional person. Written by leading experts in complementarity, duality, global optimization, and quantum computations, this collection reveals the beauty of these mathematical disciplines and investigates recent developments in global optimization, nonconvex and nonsmooth analysis, nonlinear programming, theoretical and engineering mechanics, large scale computation, quantum algorithms and computation, and information theory.

Differential equations and linear algebra are two central topics in the undergraduate mathematics curriculum. This innovative textbook allows the two subjects to be developed either separately or together, illuminating the connections between two fundamental topics, and giving increased flexibility to instructors. It can be used either as a semester-long course in differential equations, or as a one-year course in differential equations, linear algebra, and applications. Beginning with the basics of differential equations, it covers first and second order equations, graphical and numerical methods, and matrix equations. The book goes on to present the fundamentals of vector spaces, followed by eigenvalues and eigenvectors, positive definiteness, integral transform methods and applications to PDEs. The exposition illuminates the natural correspondence between solution methods for systems of equations in discrete and continuous settings. The topics draw on the physical sciences, engineering and economics, reflecting the author's distinguished career as an applied mathematician and expositor. Combining scientific computing methods and algorithms with modern data analysis techniques, including basic applications of compressive sensing and machine learning, this book develops techniques that allow for the integration of the dynamics of complex systems and big data. MATLAB is used throughout for mathematical solution strategies.

Paul Gilbert brings together an international line-up of leading scholars and researchers in the field to provide a state-of-the-art exploration of key areas in compassion research and applications. Compassion can be seen as a core element of prosocial behaviour, and explorations of the concepts and value of compassion have been extended into different aspects of life including physical and psychological therapies, schools, leadership and business. While many animals share abilities to be distress sensitive and caring of others, it is our newly evolved socially intelligent abilities that make us capable of knowingly and deliberately helping others and purposely developing skills and wisdom to do so. This book

generates many research questions whilst exploring the similarity and differences of human compassion to non-human caring and looks at how compassion changes the brain and body, affects genetic expression, manifests at a young age and is then cultivated (or not) by the social environment. *Compassion: Concepts, Research and Applications* will be essential reading for professionals, researchers and scholars interested in compassion and its applications in psychology and psychotherapy.

In this second book of what will be a four-volume series, the authors present, in a mathematically rigorous way, the essential foundations of both the theory and practice of algorithms, approximation, and optimization—essential topics in modern applied and computational mathematics. This material is the introductory framework upon which algorithm analysis, optimization, probability, statistics, machine learning, and control theory are built. This text gives a unified treatment of several topics that do not usually appear together: the theory and analysis of algorithms for mathematicians and data science students; probability and its applications; the theory and applications of approximation, including Fourier series, wavelets, and polynomial approximation; and the theory and practice of optimization, including dynamic optimization. When used in concert with the free supplemental lab materials, *Foundations of Applied Mathematics, Volume 2: Algorithms, Approximation, Optimization* teaches not only the theory but also the computational practice of modern mathematical methods. Exercises and examples build upon each other in a way that continually reinforces previous ideas, allowing students to retain learned concepts while achieving a greater depth. The mathematically rigorous lab content guides students to technical proficiency and answers the age-old question “When am I going to use this?” This textbook is geared toward advanced undergraduate and beginning graduate students in mathematics, data science, and machine learning.

Computational science is an exciting new field at the intersection of the sciences, computer science, and mathematics because much scientific investigation now involves computing as well as theory and experiment. This textbook provides students with a versatile and accessible introduction to the subject. It assumes only a background in high school algebra, enables instructors to follow tailored pathways through the material, and is the only textbook of its kind designed specifically for an introductory course in the computational science and engineering curriculum. While the text itself is generic, an accompanying website offers tutorials and files in a variety of software packages. This fully updated and expanded edition features two new chapters on agent-based simulations and modeling with matrices, ten new project modules, and an additional module on diffusion. Besides increased treatment of high-performance computing and its applications, the book also includes additional quick review questions with answers, exercises, and individual and team projects. The only introductory textbook of its kind—now fully updated and expanded Features two new chapters on agent-based simulations and modeling with matrices Increased coverage of high-performance computing and its applications Includes additional modules, review questions, exercises, and projects An online instructor's manual with exercise answers, selected project solutions, and a test bank and solutions (available only to professors) An online illustration package is available to professors Computational Science and EngineeringWellesley-Cambridge Press

Professor Ralph Kleinman was director of the Center for the Mathematics of Waves and held the UNIDEL Professorship of the University of Delaware. Before his death in 1998, he made major scientific contributions in the areas of electromagnetic scattering, wave propagation, and inverse problems. He was instrumental in bringing together the mathematic

An introduction to computational complexity theory, its connections and interactions with mathematics, and its central role in the natural and social sciences, technology, and philosophy *Mathematics and Computation* provides a broad, conceptual overview of computational complexity theory—the mathematical study of efficient computation. With important practical applications to computer science and industry, computational complexity

theory has evolved into a highly interdisciplinary field, with strong links to most mathematical areas and to a growing number of scientific endeavors. Avi Wigderson takes a sweeping survey of complexity theory, emphasizing the field's insights and challenges. He explains the ideas and motivations leading to key models, notions, and results. In particular, he looks at algorithms and complexity, computations and proofs, randomness and interaction, quantum and arithmetic computation, and cryptography and learning, all as parts of a cohesive whole with numerous cross-influences. Wigderson illustrates the immense breadth of the field, its beauty and richness, and its diverse and growing interactions with other areas of mathematics. He ends with a comprehensive look at the theory of computation, its methodology and aspirations, and the unique and fundamental ways in which it has shaped and will further shape science, technology, and society. For further reading, an extensive bibliography is provided for all topics covered. Mathematics and Computation is useful for undergraduate and graduate students in mathematics, computer science, and related fields, as well as researchers and teachers in these fields. Many parts require little background, and serve as an invitation to newcomers seeking an introduction to the theory of computation. Comprehensive coverage of computational complexity theory, and beyond High-level, intuitive exposition, which brings conceptual clarity to this central and dynamic scientific discipline Historical accounts of the evolution and motivations of central concepts and models A broad view of the theory of computation's influence on science, technology, and society Extensive bibliography The emergence of satellite technology has changed the lives of millions of people. In particular, GPS has brought an unprecedented level of accuracy to the field of geodesy. This text is a guide to the algorithms and mathematical principles that account for the success of GPS technology and replaces the authors' previous work, *Linear Algebra, Geodesy, and GPS* (1997). An initial discussion of the basic concepts, characteristics and technical aspects of different satellite systems is followed by the necessary mathematical content which is presented in a detailed and self-contained fashion. At the heart of the matter are the positioning algorithms on which GPS technology relies, the discussion of which will affirm the mathematical contents of the previous chapters. Numerous ready-to-use MATLAB codes are included for the reader. This comprehensive guide will be invaluable for engineers and academic researchers who wish to master the theory and practical application of GPS technology.

This second edition has two parts. The first part is the complete classic by Gilbert Strang and George Fix, first published in 1973. The original book demonstrates the solid mathematical foundation of the finite element idea, and the reasons for its success. The second part is a new textbook by Strang. It provides examples, codes, and exercises to connect the theory of the Finite Element Method directly to the applications. The reader will learn how to assemble the stiffness matrix  $K$  and solve the finite element equations  $KU=F$ . Discontinuous Galerkin methods with a numerical flux function are now included. Strang's approach is direct and focuses on learning finite elements by using them.

This volume in the series *Lecture Notes in Computational Science and Engineering* presents a collection of papers presented at the International Workshop on FSI, held in October 2005 in Hohenwart and organized by DFG's Research Unit 493 "FSI: Modeling, Simulation, and Optimization". The papers address partitioned and monolithic coupling approaches, methodical issues and applications, and discuss FSI from the mathematical, informatics, and engineering points of view.

Computational science and engineering (CSE) is a broad multidisciplinary and integrative area including a variety of applications in science, engineering, numerical methods, applied mathematics, and computer science disciplines. The book covers a collection of different types of applications and visions to various disciplinary key aspects, which comprises both problem-driven and methodology-driven approaches at the same time. These selected applications are:

Computational and information technologies for numerical models and large unstructured data processing Evolution of matrix computations and new concepts in computing Inverse problems covering both classical and newer approaches Integro-differential scheme (IDS) that combines finite volume and finite difference methods Smart city wireless networks Signal processing methods

When scientists describe their results or insights as 'beautiful', are they using the term differently from when they use it of a landscape, music or another person? Science and the Truthfulness of Beauty re-examines the way in which seeing beauty in the world plays the key role in scientific advances, and argues that the reliance on such a personal point of view is ultimately justified by belief that we are made in the 'image of God', as Christian and Jewish believers assert. It brings a fresh voice to the ongoing debate about faith and science, and suggests that scientists have as much explaining to do as believers when it comes to the ways they reach their conclusions.

Fundamentals of Numerical Computation is an advanced undergraduate-level introduction to the mathematics and use of algorithms for the fundamental problems of numerical computation: linear algebra, finding roots, approximating data and functions, and solving differential equations. The book is organized with simpler methods in the first half and more advanced methods in the second half, allowing use for either a single course or a sequence of two courses. The authors take readers from basic to advanced methods, illustrating them with over 200 self-contained MATLAB functions and examples designed for those with no prior MATLAB experience. Although the text provides many examples, exercises, and illustrations, the aim of the authors is not to provide a cookbook per se, but rather an exploration of the principles of cooking. The authors have developed an online resource that includes well-tested materials related to every chapter. Among these materials are lecture-related slides and videos, ideas for student projects, laboratory exercises, computational examples and scripts, and all the functions presented in the book. The book is intended for advanced undergraduates in math, applied math, engineering, or science disciplines, as well as for researchers and professionals looking for an introduction to a subject they missed or overlooked in their education.

Accessible text features over 100 reality-based examples pulled from the science, engineering, and operations research fields. Prerequisites: ordinary differential equations, continuous probability. Numerous references. Includes 27 black-and-white figures. 1978 edition.

Includes detailed step-by-step solutions to selected odd-numbered problems.

An accessible introduction to the finite element method for solving numeric problems, this volume offers the keys to an important technique in computational mathematics. Suitable for advanced undergraduate and graduate courses, it outlines clear connections with applications and considers numerous examples from a variety of science- and engineering-related specialties. This text encompasses all varieties of the basic linear partial differential equations, including elliptic, parabolic and hyperbolic problems, as well as stationary and time-dependent problems. Additional topics include finite element methods for integral equations, an introduction to nonlinear problems, and considerations of

unique developments of finite element techniques related to parabolic problems, including methods for automatic time step control. The relevant mathematics are expressed in non-technical terms whenever possible, in the interests of keeping the treatment accessible to a majority of students.

Encompasses the full range of computational science and engineering from modelling to solution, both analytical and numerical. It develops a framework for the equations and numerical methods of applied mathematics. Gilbert Strang has taught this material to thousands of engineers and scientists (and many more on MIT's OpenCourseWare 18.085-6). His experience is seen in his clear explanations, wide range of examples, and teaching method. The book is solution-based and not formula-based: it integrates analysis and algorithms and MATLAB codes to explain each topic as effectively as possible. The topics include applied linear algebra and fast solvers, differential equations with finite differences and finite elements, Fourier analysis and optimization. This book also serves as a reference for the whole community of computational scientists and engineers. Supporting resources, including MATLAB codes, problem solutions and video lectures from Gilbert Strang's 18.085 courses at MIT, are provided at [math.mit.edu/cse](http://math.mit.edu/cse).

A comprehensive treatment of wavelets for both engineers and mathematicians. The current exponential growth in graph data has forced a shift to parallel computing for executing graph algorithms. Implementing parallel graph algorithms and achieving good parallel performance have proven difficult. This book addresses these challenges by exploiting the well-known duality between a canonical representation of graphs as abstract collections of vertices and edges and a sparse adjacency matrix representation. This linear algebraic approach is widely accessible to scientists and engineers who may not be formally trained in computer science. The authors show how to leverage existing parallel matrix computation techniques and the large amount of software infrastructure that exists for these computations to implement efficient and scalable parallel graph algorithms. The benefits of this approach are reduced algorithmic complexity, ease of implementation, and improved performance.

This book gathers thousands of up-to-date equations, formulas, tables, illustrations, and explanations into one invaluable volume. It includes over a thousand pages of mathematical material as well as chapters on probability, mathematical statistics, fuzzy logic, and neural networks. It also contains computer language overviews of C, Fortran, and Pascal.

Offers students a practical knowledge of modern techniques in scientific computing.

This open access book features a selection of high-quality papers from the presentations at the International Conference on Spectral and High-Order Methods 2018, offering an overview of the depth and breadth of the activities within this important research area. The carefully reviewed papers provide a snapshot of the state of the art, while the extensive bibliography helps initiate new research directions.

The field of uncertainty quantification is evolving rapidly because of increasing emphasis on models that require quantified uncertainties for large-scale applications, novel algorithm development, and new computational architectures that facilitate implementation of these algorithms. *Uncertainty Quantification: Theory, Implementation, and Applications* provides readers with the basic concepts, theory, and algorithms necessary to quantify input and response uncertainties for simulation models arising in a broad range of disciplines. The book begins with a detailed discussion of applications where uncertainty quantification is critical for both scientific understanding and policy. It then covers concepts from probability and statistics, parameter selection techniques, frequentist and Bayesian model calibration, propagation of uncertainties, quantification of model discrepancy, surrogate model construction, and local and global sensitivity analysis. The author maintains a complementary web page where readers can find data used in the exercises and other supplementary material.

Discusses algorithms generally expressed in MATLAB for geodesy and global positioning. Three parts cover basic linear algebra, the application to the (linear and also nonlinear) science of measurement, and the GPS system and its applications. A popular article from SIAM News (June 1997) *The Mathematics of GPS* is included as an introduction. Annot

Linear algebra has become the subject to know for people in quantitative disciplines of all kinds. No longer the exclusive domain of mathematicians and engineers, it is now used everywhere there is data and everybody who works with data needs to know more. This new book from Professor Gilbert Strang, author of the acclaimed *Introduction to Linear Algebra*, now in its fifth edition, makes linear algebra accessible to everybody, not just those with a strong background in mathematics. It takes a more active start, beginning by finding independent columns of small matrices, leading to the key concepts of linear combinations and rank and column space. From there it passes on to the classical topics of solving linear equations, orthogonality, linear transformations and subspaces, all clearly explained with many examples and exercises. The last major topics are eigenvalues and the important singular value decomposition, illustrated with applications to differential equations and image compression. A final optional chapter explores the ideas behind deep learning. The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

[Copyright: 9e04029db97dccff52b9f726a2e41294](#)