1. Record Nr. UNINA9911021962903321

Autore	Calder Jeff
Titolo	Linear Algebra, Data Science, and Machine Learning / / by Jeff Calder, Peter J. Olver
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2025
ISBN	9783031937644 9783031937637
Edizione	[1st ed. 2025.]
Descrizione fisica	1 online resource (645 pages)
Collana	Springer Undergraduate Texts in Mathematics and Technology, , 1867-5514
Altri autori (Persone)	OlverPeter J
Disciplina	512.5
Soggetti	Algebras, Linear Artificial intelligence - Data processing Machine learning Mathematical optimization Probabilities Linear Algebra Data Science Machine Learning Optimization Applied Probability
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Preface -- 1 Vectors -- 2 Inner Product, Orthogonality, Norm -- 3 Matrices -- 4. How Matrices Interact with Inner Products and Norms -- 5 Eigenvalues and Singular Values -- 6 Basics of Optimization -- 7 Introduction to Machine Learning and Data -- 8 Principal Component Analysis -- 9 Graph Theory and Graph-based Learning -- 10 Neural Networks and Deep Learning -- 11 Advanced Optimization -- Bibliography -- Index.
Sommario/riassunto	This text provides a mathematically rigorous introduction to modern methods of machine learning and data analysis at the advanced undergraduate/beginning graduate level. The book is self-contained and requires minimal mathematical prerequisites. There is a strong

focus on learning how and why algorithms work, as well as developing facility with their practical applications. Apart from basic calculus, the underlying mathematics — linear algebra, optimization, elementary probability, graph theory, and statistics — is developed from scratch in a form best suited to the overall goals. In particular, the wide-ranging linear algebra components are unique in their ordering and choice of topics, emphasizing those parts of the theory and techniques that are used in contemporary machine learning and data analysis. The book will provide a firm foundation to the reader whose goal is to work on applications of machine learning and/or research into the further development of this highly active field of contemporary applied mathematics. To introduce the reader to a broad range of machine learning algorithms and how they are used in real world applications, the programming language Python is employed and offers a platform for many of the computational exercises. Python notebooks complementing various topics in the book are available on a companion GitHub site specified in the Preface, and can be easily accessed by scanning the QR codes or clicking on the links provided within the text. Exercises appear at the end of each section, including basic ones designed to test comprehension and computational skills, while others range over proofs not supplied in the text, practical computations, additional theoretical results, and further developments in the subject. The Students' Solutions Manual may be accessed from GitHub. Instructors may apply for access to the Instructors' Solutions Manual from the link supplied on the text's Springer website. The book can be used in a junior or senior level course for students majoring in mathematics with a focus on applications as well as students from other disciplines who desire to learn the tools of modern applied linear algebra and optimization. It may also be used as an introduction to fundamental techniques in data science and machine learning for advanced undergraduate and graduate students or researchers from other areas, including statistics, computer science, engineering, biology, economics and finance, and so on.