

1. Record Nr.	UNINA9910824307003321
Titolo	Nearest-neighbor methods in learning and vision : theory and practice // edited by Gregory Shakhnarovich, Trevor Darrell, Piotr Indyk
Pubbl/distr/stampa	Cambridge, Mass., : MIT Press, c2005
ISBN	1-282-09675-3 9786612096754 0-262-25695-9 1-4237-7253-9
Edizione	[1st ed.]
Descrizione fisica	1 online resource (280 p.)
Collana	Neural information processing series
Altri autori (Persone)	DarrellTrevor IndykPiotr ShakhnarovichGregory
Disciplina	006.3/1
Soggetti	Algorithms Geometry - Data processing Machine learning Nearest neighbor analysis (Statistics)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	"... held in Whistler, British Columbia ... annual conference on Neural Information Processing Systems (NIPS) in December 2003"--Pref.
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Contents; Series Foreword; Preface; 1 Introduction; I THEORY; 2 Nearest-Neighbor Searching and Metric Space Dimensions; 3 Locality-Sensitive Hashing Using Stable Distributions; II APPLICATIONS: LEARNING; 4 New Algorithms for Efficient High-Dimensional Nonparametric Classification; 5 Approximate Nearest Neighbor Regression in Very High Dimensions; 6 Learning Embeddings for Fast Approximate Nearest Neighbor Retrieval; III APPLICATIONS: VISION; 7 Parameter-Sensitive Hashing for Fast Pose Estimation; 8 Contour Matching Using Approximate Earth Mover's Distance 9 Adaptive Mean Shift Based Clustering in High Dimensions 10 Object Recognition using Locality Sensitive Hashing of Shape Contexts; Contributors; Index
Sommario/riassunto	Regression and classification methods based on similarity of the input to stored examples have not been widely used in applications involving

very large sets of high-dimensional data. Recent advances in computational geometry and machine learning, however, may alleviate the problems in using these methods on large data sets. This volume presents theoretical and practical discussions of nearest-neighbor (NN) methods in machine learning and examines computer vision as an application domain in which the benefit of these advanced methods is often dramatic. It brings together contributions from researchers in theory of computation, machine learning, and computer vision with the goals of bridging the gaps between disciplines and presenting state-of-the-art methods for emerging applications. The contributors focus on the importance of designing algorithms for NN search, and for the related classification, regression, and retrieval tasks, that remain efficient even as the number of points or the dimensionality of the data grows very large. The book begins with two theoretical chapters on computational geometry and then explores ways to make the NN approach practicable in machine learning applications where the dimensionality of the data and the size of the data sets make the naive methods for NN search prohibitively expensive. The final chapters describe successful applications of an NN algorithm, locality-sensitive hashing (LSH), to vision tasks.
