1.	Record Nr.	UNINA9910276947203321
	Autore	Goodfellow Ian
	Titolo	Deep learning / / Ian Goodfellow, Yoshua Bengio and Aaron Courville
	Pubbl/distr/stampa	Cambridge, Massachusetts ; ; London, England : , : The MIT Press, , [2016] ©2016
	ISBN	0-262-33737-1
	Descrizione fisica	1 online resource (xxii, 775 pages) : illustrations
	Collana	Adaptive computation and machine learning
	Disciplina	006.31
	Soggetti	Machine learning
	Lingua di pubblicazione	Inglese
	Formato	Materiale a stampa
	Livello bibliografico	Monografia
	Nota di bibliografia	Includes bibliographical references (pages 711-766) and index.
	Nota di contenuto	Introduction -- Applied math and machine learning basics. Linear algebra -- Probability and information theory -- Numerical computation -- Machine learning basics -- Deep networks: modern practices. Deep feedforward networks -- Regularization for deep learning -- Optimization for training deep models -- Convolutional networks -- Sequence modeling: recurrent and recursive nets -- Practical methodology -- Applications -- Deep learning research. Linear factor models -- Autoencoders -- Representation learning -- Structured probabilistic models for deep learning -- Monte Carlo methods -- Confronting the partition function -- Approximate inference -- Deep generative models.
	Sommario/riassunto	An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts

in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.