

1. Record Nr.	UNINA9910861088903321
Autore	Aggarwal Charu C
Titolo	Probability and Statistics for Machine Learning : A Textbook / / by Charu C. Aggarwal
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2024
ISBN	3-031-53282-1
Edizione	[1st ed. 2024.]
Descrizione fisica	1 online resource (530 pages)
Disciplina	006.31
Soggetti	Machine learning Machine Learning
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Chapter. 1. Probability and Statistics: An Introduction -- Chapter. 2. Summarizing and Visualizing Data -- Chapter. 3. Probability Basics and Random Variables -- Chapter. 4. Probability Distributions -- Chapter. 5. Hypothesis Testing and Confidence Intervals -- Chapter. 6. Reconstructing Probability Distributions from Data -- Chapter. 7. Regression -- Chapter. 8. Classification: A Probabilistic View -- Chapter. 9. Unsupervised Learning: A Probabilistic View -- Chapter. 10. Discrete State Markov Processes -- Chapter. 11. Probabilistic Inequalities and Extreme Value Analysis -- Bibliography -- Index.
Sommario/riassunto	This book covers probability and statistics from the machine learning perspective. The chapters of this book belong to three categories: 1. The basics of probability and statistics: These chapters focus on the basics of probability and statistics, and cover the key principles of these topics. Chapter 1 provides an overview of the area of probability and statistics as well as its relationship to machine learning. The fundamentals of probability and statistics are covered in Chapters 2 through 5. 2. From probability to machine learning: Many machine learning applications are addressed using probabilistic models, whose parameters are then learned in a data-driven manner. Chapters 6 through 9 explore how different models from probability and statistics are applied to machine learning. Perhaps the most important tool that bridges the gap from data to probability is maximum-likelihood estimation, which is a foundational concept from the perspective of

machine learning. This concept is explored repeatedly in these chapters. 3. Advanced topics: Chapter 10 is devoted to discrete-state Markov processes. It explores the application of probability and statistics to a temporal and sequential setting, although the applications extend to more complex settings such as graphical data. Chapter 11 covers a number of probabilistic inequalities and approximations. The style of writing promotes the learning of probability and statistics simultaneously with a probabilistic perspective on the modeling of machine learning applications. The book contains over 200 worked examples in order to elucidate key concepts. Exercises are included both within the text of the chapters and at the end of the chapters. The book is written for a broad audience, including graduate students, researchers, and practitioners.
