


The book was found

Introduction To Machine Learning (Adaptive Computation And Machine Learning Series)



Synopsis

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing. Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods.

Book Information

File Size: 15972 KB

Print Length: 640 pages

Publisher: The MIT Press; 3 edition (August 22, 2014)

Publication Date: August 22, 2014

Sold by: Digital Services LLC

Language: English

ASIN: B00NLVNGLA

Text-to-Speech: Not enabled

X-Ray: Not Enabled

Word Wise: Not Enabled

Lending: Not Enabled

Enhanced Typesetting: Not Enabled

Best Sellers Rank: #534,839 Paid in Kindle Store (See Top 100 Paid in Kindle Store) #120

inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #556 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics #8786 inÂ Kindle Store > Kindle eBooks > Computers & Technology

Customer Reviews

I understand ML very well, and I find this text nearly impossible to penetrate. Formulas are reduced to their most rudimentary forms. Sure it is impressive that the author obviously has a good grasp on the topic, but there are virtually no explanations behind the math. This book was written just to show off, not to teach. Definitely the most pompous book on ML I've ever seen.

Very nice and clear content. Thanks....

Item as described. Fast shipment.

good for beginners

[Download to continue reading...](#)

Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series) Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and Machine Learning) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear

Regression and Logistic Regression in Python (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms Introduction to Automata Theory, Languages, and Computation (3rd Edition) Introduction to Automata Theory, Languages, and Computation Introduction to Automata Theory, Languages, and Computation (2nd Edition)

[Dmca](#)