

Learning From Data Yaser Mostafa

A comprehensive overview of data mining from an algorithmic perspective, integrating related concepts from machine learning and statistics.

Big Data Analytics with Spark is a step-by-step guide for learning Spark, which is an open-source fast and general-purpose cluster computing framework for large-scale data analysis. You will learn how to use Spark for different types of big data analytics projects, including batch, interactive, graph, and stream data analysis as well as machine learning. In addition, this book will help you become a much sought-after Spark expert. Spark is one of the hottest Big Data technologies. The amount of data generated today by devices, applications and users is

Online Library Learning From Data Yaser Mostafa

exploding. Therefore, there is a critical need for tools that can analyze large-scale data and unlock value from it. Spark is a powerful technology that meets that need. You can, for example, use Spark to perform low latency computations through the use of efficient caching and iterative algorithms; leverage the features of its shell for easy and interactive Data analysis; employ its fast batch processing and low latency features to process your real time data streams and so on. As a result, adoption of Spark is rapidly growing and is replacing Hadoop MapReduce as the technology of choice for big data analytics. This book provides an introduction to Spark and related big-data technologies. It covers Spark core and its add-on libraries, including Spark SQL, Spark Streaming, GraphX, and MLlib. Big Data Analytics with Spark is therefore written for busy professionals

Online Library Learning From Data Yaser Mostafa

who prefer learning a new technology from a consolidated source instead of spending countless hours on the Internet trying to pick bits and pieces from different sources. The book also provides a chapter on Scala, the hottest functional programming language, and the program that underlies Spark. You'll learn the basics of functional programming in Scala, so that you can write Spark applications in it. What's more, Big Data Analytics with Spark provides an introduction to other big data technologies that are commonly used along with Spark, like Hive, Avro, Kafka and so on. So the book is self-sufficient; all the technologies that you need to know to use Spark are covered. The only thing that you are expected to know is programming in any language. There is a critical shortage of people with big data expertise, so companies are willing to pay top dollar

Online Library Learning From Data Yaser Mostafa

for people with skills in areas like Spark and Scala. So reading this book and absorbing its principles will provide a boost—possibly a big boost—to your career. An accessible introduction and essential reference for an approach to machine learning that creates highly accurate prediction rules by combining many weak and inaccurate ones. Boosting is an approach to machine learning based on the idea of creating a highly accurate predictor by combining many weak and inaccurate “rules of thumb.” A remarkably rich theory has evolved around boosting, with connections to a range of topics, including statistics, game theory, convex optimization, and information geometry. Boosting algorithms have also enjoyed practical success in such fields as biology, vision, and speech processing. At various times in its history, boosting has been perceived as

Online Library Learning From Data Yaser Mostafa

mysterious, controversial, even paradoxical. This book, written by the inventors of the method, brings together, organizes, simplifies, and substantially extends two decades of research on boosting, presenting both theory and applications in a way that is accessible to readers from diverse backgrounds while also providing an authoritative reference for advanced researchers. With its introductory treatment of all material and its inclusion of exercises in every chapter, the book is appropriate for course use as well. The book begins with a general introduction to machine learning algorithms and their analysis; then explores the core theory of boosting, especially its ability to generalize; examines some of the myriad other theoretical viewpoints that help to explain and understand boosting; provides practical extensions of boosting for more

Online Library Learning From Data Yaser Mostafa

complex learning problems; and finally presents a number of advanced theoretical topics. Numerous applications and practical illustrations are offered throughout.

November 28-December 1, 1994, Denver, Colorado NIPS is the longest running annual meeting devoted to Neural Information Processing Systems. Drawing on such disparate domains as neuroscience, cognitive science, computer science, statistics, mathematics, engineering, and theoretical physics, the papers collected in the proceedings of NIPS7 reflect the enduring scientific and practical merit of a broad-based, inclusive approach to neural information processing. The primary focus remains the study of a wide variety of learning algorithms and architectures, for both supervised and unsupervised learning. The 139 contributions are divided into eight

Online Library Learning From Data Yaser Mostafa

parts: Cognitive Science, Neuroscience, Learning Theory, Algorithms and Architectures, Implementations, Speech and Signal Processing, Visual Processing, and Applications. Topics of special interest include the analysis of recurrent nets, connections to HMMs and the EM procedure, and reinforcement- learning algorithms and the relation to dynamic programming. On the theoretical front, progress is reported in the theory of generalization, regularization, combining multiple models, and active learning. Neuroscientific studies range from the large-scale systems such as visual cortex to single-cell electrotonic structure, and work in cognitive scientific is closely tied to underlying neural constraints. There are also many novel applications such as tokamak plasma control, Glove-Talk, and hand tracking, and a variety of hardware implementations, with particular focus on

Online Library Learning From Data Yaser Mostafa

analog VLSI.

A Short Course

Volume 1: Imaging and Signal Analysis

Data Mining and Analysis

*Database Design Using Entity-
Relationship Diagrams*

A Set of Lectures

From Big Data to Smart Data

Pyomo – Optimization Modeling in Python

The annual conference on Neural Information Processing Systems (NIPS) is the flagship conference on neural computation. It draws preeminent academic researchers from around the world and is widely considered to be a showcase conference for new developments in network algorithms and architectures. The broad range of interdisciplinary research areas represented includes computer

Online Library Learning From Data Yaser Mostafa

science, neuroscience, statistics, physics, cognitive science, and many branches of engineering, including signal processing and control theory. Only about 30 percent of the papers submitted are accepted for presentation at NIPS, so the quality is exceptionally high. These proceedings contain all of the papers that were presented. A state-of-the-art guide for the implementation of distributed simulation technology. The rapid expansion of the Internet and commodity parallel computers has made parallel and distributed simulation (PADS) a hot technology indeed. Applications abound not only in the analysis of complex systems such as transportation or

Online Library Learning From Data Yaser Mostafa

the next-generation Internet, but also in computer-generated virtual worlds for military and professional training, interactive computer games, and the entertainment industry. In this book, PADS expert Richard M. Fujimoto provides software developers with cutting-edge techniques for speeding up the execution of simulations across multiple processors and dealing with data distribution over wide area networks ,including the Internet. With an emphasis on parallel and distributed discrete event simulation technologies, Dr. Fujimoto compiles and consolidates research results in the field spanning the last twenty years,

Online Library Learning From Data Yaser Mostafa

discussing the use of parallel and distributed computers in both the modeling and analysis of system behavior and the creation of distributed virtual environments. While other books on PADS concentrate on applications, *Parallel and Distributed Simulation Systems* clearly shows how to implement the technology. It explains in detail the synchronization algorithms needed to properly realize the simulations, including an in-depth discussion of time warp and advanced optimistic techniques. Finally, the book is richly supplemented with references, tables and illustrations, and examples of contemporary systems such as the Department of

Online Library Learning From Data Yaser Mostafa

Defense's High Level Architecture (HLA), which has become the standard architecture for defense programs in the United States. Neural networks are a family of powerful machine learning models and this book focuses on their application to natural language data. The first half of the book (Parts I and II) covers the basics of supervised machine learning and feed-forward neural networks, the basics of working with machine learning over language data, and the use of vector-based rather than symbolic representations for words. It also covers the computation-graph abstraction, which allows to easily define and train arbitrary neural networks,

Online Library Learning From Data Yaser Mostafa

and is the basis behind the design of contemporary neural network software libraries. The second part of the book (Parts III and IV) introduces more specialized neural network architectures, including 1D convolutional neural networks, recurrent neural networks, conditioned-generation models, and attention-based models. These architectures and techniques are the driving force behind state-of-the-art algorithms for machine translation, syntactic parsing, and many other applications. Finally, we also discuss tree-shaped networks, structured prediction, and the prospects of multi-task learning.

Neural Engineering for Autism

Online Library Learning From Data Yaser Mostafa

Spectrum Disorder, Volume One: Imaging and Signal Analysis Techniques presents the latest advances in neural engineering and biomedical engineering as applied to the clinical diagnosis and treatment of Autism Spectrum Disorder (ASD). Advances in the role of neuroimaging, infrared spectroscopy, sMRI, fMRI, DTI, social behaviors and suitable data analytics useful for clinical diagnosis and research applications for Autism Spectrum Disorder are covered, including relevant case studies. The application of brain signal evaluation, EEG analytics, feature selection, and analysis of blood oxygen level-dependent (BOLD) signals are presented for

Online Library Learning From Data Yaser Mostafa

detection and estimation of the degree of ASD. Presents applications of Neural Engineering and other Machine Learning techniques for the diagnosis of Autism Spectrum Disorder (ASD) Includes in-depth technical coverage of imaging and signal analysis techniques, including coverage of functional MRI, neuroimaging, infrared spectroscopy, sMRI, fMRI, DTI, and neuroanatomy of autism Covers Signal Analysis for the detection and estimation of Autism Spectrum Disorder (ASD), including brain signal analysis, EEG analytics, feature selection, and analysis of blood oxygen level-dependent (BOLD) signals for ASD Written to

Online Library Learning From Data Yaser Mostafa

help engineers, computer scientists, researchers and clinicians understand the technology and applications of Neural Engineering for the detection and diagnosis of Autism Spectrum Disorder (ASD)

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

The Lion Way

Machine Learning Refined

Introduction to Statistical Machine Learning

Machine Learning

Deep Learning for NLP and Speech Recognition

Portfolio Management under Stress

Machine learning (ML) is changing virtually every aspect of our lives.

Online Library Learning From Data Yaser Mostafa

Today ML algorithms accomplish tasks that until recently only expert humans could perform. As it relates to finance, this is the most exciting time to adopt a disruptive technology that will transform how everyone invests for generations. Readers will learn how to structure Big data in a way that is amenable to ML algorithms; how to conduct research with ML algorithms on that data; how to use supercomputing methods; how to backtest your discoveries while avoiding false positives. The book addresses real-life problems faced by practitioners on a daily basis, and explains scientifically sound solutions using math, supported by code and examples. Readers become active users who can test the proposed solutions in their

Online Library Learning From Data Yaser Mostafa

particular setting. Written by a recognized expert and portfolio manager, this book will equip investment professionals with the groundbreaking tools needed to succeed in modern finance.

STATISTICS: LEARNING FROM DATA, Second Edition, helps you learn to think like a statistician. It pays particular attention to areas that students often struggle with -- probability, hypothesis testing, and selecting an appropriate method of analysis. Supported by learning objectives, real-data examples and exercises, and technology notes, this book helps you to develop conceptual understanding, mechanical proficiency, and the ability to put knowledge into practice. Important Notice: Media content referenced within the

Online Library Learning From Data Yaser Mostafa

product description or the product text may not be available in the ebook version.

Entity-relationship (E-R) diagrams are time-tested models for database development well-known for their usefulness in mapping out clear database designs. Also commonly known is how difficult it is to master them. With this comprehensive guide, database designers and developers can quickly learn all the ins and outs of E-R diagramming to become experts. Presenting a theory of the theoryless, a computer scientist provides a model of how effective behavior can be learned even in a world as complex as our own, shedding new light on human nature.

Information Theory, Inference and

Online Library Learning From Data Yaser Mostafa

**Learning Algorithms
Learning Plus Intelligent
Optimization
Nature's Algorithms for Learning
and Prospering in a Complex World
Foundations and Algorithms
Foundations, Algorithms, and
Applications**

**On the Path to AI
*A pragmatic approach to
Big Data by taking the
reader on a journey
between Big Data (what it
is) and the Smart Data
(what it is for). Today's
decision making can be
reached via information
(related to the data),
knowledge (related to***

people and processes), and timing (the capacity to decide, act and react at the right time). The huge increase in volume of data traffic, and its format (unstructured data such as blogs, logs, and video) generated by the “digitalization” of our world modifies radically our relationship to the space (in motion) and time, dimension and by capillarity, the enterprise vision of performance monitoring and optimization.

This text is a semester

course in the basic mathematical and theoretical foundations of computer science. Students who make heavy use of computing should learn these foundations well, setting a base for a follow-on course in algorithms. A solid theoretical and algorithmic foundation in computer science sets the stage for developing good programs, programs that work, always and efficiently. Each chapter is a lecture that has been taught as such. Part I

starts with basic logic, proofs and discrete mathematics, including: induction, recursion, summation, asymptotics and number theory. We then continue with graphs, counting and combinatorics, and wrap up the coverage of discrete mathematics with discrete probability. Part II presents the blockbuster application of discrete mathematics: the digital computer and a theory of computing. The goal is to understand what a computer can and

cannot do. We start small, with automata, and end big with Turing Machines. Our approach is Socratic. The reader is encouraged to participate actively in the learning process by doing the quizzes and exercises that are liberally sprinkled through the text. The pace and level is appropriate for readers with one year of training in programming and calculus (college sophomores). This in-depth guide provides managers with a

solid understanding of data and data trends, the opportunities that it can offer to businesses, and the dangers of these technologies. Written in an accessible style, Steven Finlay provides a contextual roadmap for developing solutions that deliver benefits to organizations.

Table of contents

Advances in Neural Information Processing Systems 12

Neural Network Methods in Natural Language Processing

***Discrete Mathematics and
Computing***

***Advances in Neural
Information Processing
Systems 10***

***Data Wrangling with
Pandas, NumPy, and
IPython***

***Probably Approximately
Correct***

Learning to Learn

**Providing a unique approach
to machine learning, this
text contains fresh and
intuitive, yet rigorous,
descriptions of all
fundamental concepts
necessary to conduct
research, build products,**

tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational

Online Library Learning From Data Yaser Mostafa

exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

This up-to-date business analytics textbook (published in July 2020) will get you harnessing the power of the R programming language to: manipulate and

Online Library Learning From Data Yaser Mostafa

model data, discover and communicate insight, to visually communicate that insight, and successfully advocate for change within an organization. Book Description A frequent teaching-award winning professor with an analytics-industry background shares his hands-on guide to learning business analytics. It is the first textbook addressing a complete and modern business analytics workflow that includes data manipulation, data visualization, modelling business problems with

graphical models, translating graphical models into code, and presenting insights back to stakeholders. Book Highlights Content that is accessible to anyone, even most analytics beginners. If you have taken a stats course, you are good to go. Assumes no knowledge of the R programming language. Provides introduction to R, RStudio, and the Tidyverse. Provides a solid foundation and an implementable workflow for anyone wading into the Bayesian inference waters.

Online Library Learning From Data Yaser Mostafa

Provides a complete workflow within the R-ecosystem; there is no need to learn several programming languages or work through clunky interfaces between software tools. First book introducing two powerful R-packages - `causact` for visual modelling of business problems and `greta` which is an R interface to `TensorFlow` used for Bayesian inference. Uses the intuitive coding practices of the `tidyverse` including using `dplyr` for data manipulation and `ggplot2`

for data visualization.

Datasets that are freely and easily accessible. Code for generating all results and almost every visualization used in the textbook. Do not learn statistical computation or fancy math in a vacuum, learn it through this guide within the context of solving business problems.

This open access book explores machine learning and its impact on how we make sense of the world. It does so by bringing together two 'revolutions' in a surprising analogy: the revolution of machine

learning, which has placed computing on the path to artificial intelligence, and the revolution in thinking about the law that was spurred by Oliver Wendell Holmes Jr in the last two decades of the 19th century. Holmes reconceived law as prophecy based on experience, prefiguring the buzzwords of the machine learning age-prediction based on datasets. On the path to AI introduces readers to the key concepts of machine learning, discusses the potential applications and limitations

Online Library Learning From Data Yaser Mostafa

of predictions generated by machines using data, and informs current debates amongst scholars, lawyers and policy makers on how it should be used and regulated wisely.

Technologists will also find useful lessons learned from the last 120 years of legal grappling with accountability, explainability, and biased data.

Over the past three decades or so, research on machine learning and data mining has led to a wide variety of algorithms that learn

general functions from experience. As machine learning is maturing, it has begun to make the successful transition from academic research to various practical applications. Generic techniques such as decision trees and artificial neural networks, for example, are now being used in various commercial and industrial applications. Learning to Learn is an exciting new research direction within machine learning. Similar to traditional machine-learning algorithms, the methods

described in Learning to Learn induce general functions from experience. However, the book investigates algorithms that can change the way they generalize, i.e., practice the task of learning itself, and improve on it. To illustrate the utility of learning to learn, it is worthwhile comparing machine learning with human learning. Humans encounter a continual stream of learning tasks. They do not just learn concepts or motor skills, they also learn bias, i.e., they learn how to

generalize. As a result, humans are often able to generalize correctly from extremely few examples - often just a single example suffices to teach us a new thing. A deeper understanding of computer programs that improve their ability to learn can have a large practical impact on the field of machine learning and beyond. In recent years, the field has made significant progress towards a theory of learning to learn along with practical new algorithms, some of which led to impressive results in

real-world applications. Learning to Learn provides a survey of some of the most exciting new research approaches, written by leading researchers in the field. Its objective is to investigate the utility and feasibility of computer programs that can learn how to learn, both from a practical and a theoretical point of view.

A Bayesian-Net Approach to Coherent Asset Allocation

A Mathematical Approach

Myths, Misconceptions and Methods

Intro to Bayesian Business

**Analytics in the R Ecosystem
Advances in Neural
Information Processing
Systems 9
Neural Engineering
Techniques for Autism
Spectrum Disorder
Fundamental Concepts and
Algorithms**

**Covering all the main
approaches in state-of-the-art
machine learning research,
this will set a new standard as
an introductory textbook.**

**129 6.2 Representation of
hints. 131 6.3 Monotonicity
hints .. . 134 6.4 Theory .. .
. 139 6.4.1 Capacity results
140 6.4.2 Decision boundaries
144 6.5 Conclusion 145 6.6**

References.....	... 146	7
Analysis and Synthesis Tools for Robust SPRness	147	C.
Mosquera, J.R. Hernandez, F. Perez-Gonzalez	7.1	
Introduction.....	147	7.2
SPR Analysis of Uncertain Systems.	153	7.2.1
The Polytopic Case .	155	7.2.2
The ZP-Ball Case	157	7.2.3
The Roots Space Case ...	159	7.3
Synthesis of LTI Filters for Robust SPR Problems	161	
7.3.1 Algebraic Design for Two Plants	161	7.3.2
Algebraic Design for Three or More Plants	164	7.3.3
Approximate Design Methods.	165	7.4
Experimental results	167	7.5
Conclusions	168	7.6
References	169	8
Boundary Methods for		

Online Library Learning From Data Yaser Mostafa

Distribution Analysis	173
J.L. Sancho et al.	8.1 Introduction
.....	173
8.1.1 Building a Classifier System	175
8.2 Motivation	176
8.3 Boundary Methods as Feature-Set Evaluation	177
8.3.1 Results	179
8.3.2 Feature Set Evaluation using Boundary Methods: Summary	182

..

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. “Written by three experts in the field, Deep Learning is the only comprehensive book on

the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX

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

Online Library Learning From Data Yaser Mostafa

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

Online Library Learning From Data Yaser Mostafa

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

Online Library Learning From Data Yaser Mostafa

both readers and instructors.

Daily Coding Problem

contains a wide variety of questions inspired by real programming interviews, with in-depth solutions that clearly take you through each core concept. You'll learn about: *

Linked Lists * Arrays * Heaps

*** Trees * Graphs ***

Randomized Algorithms *

Backtracking * Dynamic

Programming * Stacks and

Queues * Bit Manipulation *

System Design

An Algorithmic Perspective

A Practitioner's Guide to

Using Spark for Large Scale

Data Analysis

The Art and Science of

Algorithms that Make Sense

of Data

Online Library Learning From Data Yaser Mostafa

Learning from Data Statistics: Learning from Data Law's Prophecies and the Conceptual Foundations of the Machine Learning Age Intelligent Methods in Signal Processing and Communications

This book provides a complete and comprehensive reference/guide to Pyomo (Python Optimization Modeling Objects) for both beginning and advanced modelers, including students at the undergraduate and graduate levels, academic researchers, and practitioners. The text illustrates the breadth of the modeling and analysis capabilities that are supported

Online Library Learning From Data Yaser Mostafa

by the software and support of complex real-world applications. Pyomo is an open source software package for formulating and solving large-scale optimization and operations research problems. The text begins with a tutorial on simple linear and integer programming models. A detailed reference of Pyomo's modeling components is illustrated with extensive examples, including a discussion of how to load data from data sources like spreadsheets and databases. Chapters describing advanced modeling capabilities for nonlinear and stochastic

Online Library Learning From Data Yaser Mostafa

optimization are also included. The Pyomo software provides familiar modeling features within Python, a powerful dynamic programming language that has a very clear, readable syntax and intuitive object orientation. Pyomo includes Python classes for defining sparse sets, parameters, and variables, which can be used to formulate algebraic expressions that define objectives and constraints. Moreover, Pyomo can be used from a command-line interface and within Python's interactive command environment, which makes it easy to create Pyomo models,

Online Library Learning From Data Yaser Mostafa

apply a variety of optimizers, and examine solutions. The software supports a different modeling approach than commercial AML (Algebraic Modeling Languages) tools, and is designed for flexibility, extensibility, portability, and maintainability but also maintains the central ideas in modern AMLs.

Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural

Online Library Learning From Data Yaser Mostafa

networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from

Online Library Learning From Data Yaser Mostafa

scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks

Online Library Learning From Data Yaser Mostafa

Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep

Online Library Learning From Data Yaser Mostafa

learning: why you should learn
it Fundamental concepts: how
do machines learn?

Introduction to neural
prediction: forward

propagation Introduction to
neural learning: gradient

descent Learning multiple
weights at a time: generalizing
gradient descent Building your
first deep neural network:

introduction to

backpropagation How to

picture neural networks: in
your head and on paper

Learning signal and ignoring
noise: introduction to

regularization and batching

Modeling probabilities and
nonlinearities: activation

Online Library Learning From Data Yaser Mostafa

functions Neural learning about
edges and corners: intro to
convolutional neural networks
Neural networks that
understand language: king -
man + woman == ? Neural
networks that write like
Shakespeare: recurrent layers
for variable-length data
Introducing automatic
optimization: let's build a deep
learning framework Learning to
write like Shakespeare: long
short-term memory Deep
learning on unseen data:
introducing federated learning
Where to go from here: a brief
guide
Traditional books on machine
learning can be divided into

Online Library Learning From Data Yaser Mostafa

two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but This textbook explains Deep Learning Architecture, with applications to various NLP Tasks, including Document Classification, Machine Translation, Language Modeling, and Speech Recognition. With the widespread adoption of deep

Online Library Learning From Data Yaser Mostafa

learning, natural language processing (NLP), and speech applications in many areas (including Finance, Healthcare, and Government) there is a growing need for one comprehensive resource that maps deep learning techniques to NLP and speech and provides insights into using the tools and libraries for real-world applications. Deep Learning for NLP and Speech Recognition explains recent deep learning methods applicable to NLP and speech, provides state-of-the-art approaches, and offers real-world case studies with code to provide hands-on experience.

Online Library Learning From Data Yaser Mostafa

Many books focus on deep learning theory or deep learning for NLP-specific tasks while others are cookbooks for tools and libraries, but the constant flux of new algorithms, tools, frameworks, and libraries in a rapidly evolving landscape means that there are few available texts that offer the material in this book. The book is organized into three parts, aligning to different groups of readers and their expertise. The three parts are: Machine Learning, NLP, and Speech Introduction The first part has three chapters that introduce readers to the fields of NLP, speech

Online Library Learning From Data Yaser Mostafa

recognition, deep learning and machine learning with basic theory and hands-on case studies using Python-based tools and libraries. Deep Learning Basics The five chapters in the second part introduce deep learning and various topics that are crucial for speech and text processing, including word embeddings, convolutional neural networks, recurrent neural networks and speech recognition basics. Theory, practical tips, state-of-the-art methods, experimentations and analysis in using the methods discussed in theory on real-world tasks. Advanced Deep Learning

Online Library Learning From Data Yaser Mostafa

Techniques for Text and
Speech The third part has five
chapters that discuss the latest
and cutting-edge research in
the areas of deep learning that
intersect with NLP and speech.
Topics including attention
mechanisms, memory
augmented networks, transfer
learning, multi-task learning,
domain adaptation,
reinforcement learning, and
end-to-end deep learning for
speech recognition are covered
using case studies.

Proceedings of the 1999
Conference

Daily Coding Problem

Digital Enterprise Design &
Management

Online Library Learning From Data Yaser Mostafa

Proceedings of the 1997
Conference
Concepts, Tools, and
Techniques to Build Intelligent
Systems
Predictive Analytics, Data
Mining and Big Data
Advances in Neural Information
Processing Systems 7
Through a series of recent
breakthroughs, deep learning has
boosted the entire field of machine
learning. Now, even programmers
who know close to nothing about
this technology can use simple,
efficient tools to implement
programs capable of learning from
data. This practical book shows you
how. By using concrete examples,
minimal theory, and two production-

Online Library Learning From Data Yaser Mostafa

ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You ' ll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you ' ve learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees,

Online Library Learning From Data Yaser Mostafa

random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

This open access book comprehensively covers the fundamentals of clinical data science, focusing on data collection, modelling and clinical applications. Topics covered in the first section on data collection include: data sources, data at scale (big data), data stewardship (FAIR data) and related privacy concerns. Aspects of predictive modelling

Online Library Learning From Data Yaser Mostafa

using techniques such as classification, regression or clustering, and prediction model validation will be covered in the second section. The third section covers aspects of (mobile) clinical decision support systems, operational excellence and value-based healthcare. Fundamentals of Clinical Data Science is an essential resource for healthcare professionals and IT consultants intending to develop and refine their skills in personalized medicine, using solutions based on large datasets from electronic health records or telemonitoring programmes. The book 's promise is "no math, no code" and will explain the topics in a style that is

Online Library Learning From Data Yaser Mostafa

optimized for a healthcare audience.

Learning and Intelligent Optimization (LION) is the combination of learning from data and optimization applied to solve complex and dynamic problems. The LION way is about increasing the automation level and connecting data directly to decisions and actions. More power is directly in the hands of decision makers in a self-service manner, without resorting to intermediate layers of data scientists. LION is a complex array of mechanisms, like the engine in an automobile, but the user (driver) does not need to know the inner workings of the engine in order to realize its tremendous

Online Library Learning From Data Yaser Mostafa

benefits. LION's adoption will create a prairie fire of innovation which will reach most businesses in the next decades. Businesses, like plants in wildfire-prone ecosystems, will survive and prosper by adapting and embracing LION techniques, or they risk being transformed from giant trees to ashes by the spreading competition.

Get complete instructions for manipulating, processing, cleaning, and crunching datasets in Python. Updated for Python 3.6, the second edition of this hands-on guide is packed with practical case studies that show you how to solve a broad set of data analysis problems effectively. You ' ll learn the latest versions of pandas, NumPy,

Online Library Learning From Data Yaser Mostafa

IPython, and Jupyter in the process. Written by Wes McKinney, the creator of the Python pandas project, this book is a practical, modern introduction to data science tools in Python. It ' s ideal for analysts new to Python and for Python programmers new to data science and scientific computing. Data files and related material are available on GitHub. Use the IPython shell and Jupyter notebook for exploratory computing Learn basic and advanced features in NumPy (Numerical Python) Get started with data analysis tools in the pandas library Use flexible tools to load, clean, transform, merge, and reshape data Create informative visualizations with

Online Library Learning From Data Yaser Mostafa

matplotlib Apply the pandas
groupby facility to slice, dice, and
summarize datasets Analyze and
manipulate regular and irregular
time series data Learn how to solve
real-world data analysis problems
with thorough, detailed examples
Get Exceptionally Good at Coding
Interviews by Solving One Problem
Every Day
Boosting
Grokking Deep Learning
Fundamentals of Clinical Data
Science
A Business Analyst's Introduction to
Business Analytics
Deep Learning Architectures
Deep Learning
Machine learning allows
computers to learn and discern

Online Library Learning From Data Yaser Mostafa

patterns without actually being programmed. When Statistical techniques and machine learning are combined together they are a powerful tool for analysing various kinds of data in many computer science/engineering areas including, image processing, speech processing, natural language processing, robot control, as well as in fundamental sciences such as biology, medicine, astronomy, physics, and materials. Introduction to Statistical Machine Learning provides a general introduction to machine learning that covers a wide

Online Library Learning From Data Yaser Mostafa

range of topics concisely and will help you bridge the gap between theory and practice. Part I discusses the fundamental concepts of statistics and probability that are used in describing machine learning algorithms. Part II and Part III explain the two major approaches of machine learning techniques; generative methods and discriminative methods. While Part III provides an in-depth look at advanced topics that play essential roles in making machine learning algorithms more useful in practice. The accompanying MATLAB/Octave programs

Online Library Learning From Data Yaser Mostafa

provide you with the necessary practical skills needed to accomplish a wide range of data analysis tasks. Provides the necessary background material to understand machine learning such as statistics, probability, linear algebra, and calculus. Complete coverage of the generative approach to statistical pattern recognition and the discriminative approach to statistical machine learning. Includes MATLAB/Octave programs so that readers can test the algorithms numerically and acquire both mathematical and practical skills in a wide range of data analysis tasks

Online Library Learning From Data Yaser Mostafa

Discusses a wide range of applications in machine learning and statistics and provides examples drawn from image processing, speech processing, natural language processing, robot control, as well as biology, medicine, astronomy, physics, and materials.

A rigorous presentation of a novel methodology for asset allocation in financial portfolios under conditions of market distress.

The annual conference on Neural Information Processing Systems (NIPS) is the flagship conference on neural

Online Library Learning From Data Yaser Mostafa

computation. These proceedings contain all of the papers that were presented. This book contains all refereed papers that were accepted to the second edition of the « Digital Enterprise Design & Management » (DED&M 2014) international conference that took place in Paris (France) from February 4 to February 5, 2014. These proceedings cover the most recent trends in the emerging field of Digital Enterprise, both from an academic and a professional perspective. A special focus is put on digital uses, digital strategies, digital infrastructures

Online Library Learning From Data Yaser Mostafa

and digital governance from an Enterprise Architecture point of view. The DED&M 2014 conference is organized under the guidance of the Center of Excellence on Systems Architecture, Management, Economy and Strategy and benefits from the supports of both the Orange – Ecole Polytechnique – Télécom ParisTech “ Innovation and Regulation ” Chair and the Dassault Aviation – DCNS – DGA – Thales – Ecole Polytechnique – ENSTA ParisTech – Télécom ParisTech “ Complex Systems Engineering ” Chair.

Online Library Learning From Data Yaser Mostafa

Advances in Financial Machine
Learning

Data Structures Using C++

Python for Data Analysis

Proceedings of the Second

International Conference on

Digital Enterprise Design and

Management DED&M 2014

Proceedings of the 1996

Conference

Parallel and Distributed

Simulation Systems

Big Data Analytics with Spark

The annual conference on

Neural Information

Processing Systems (NIPS) is

the flagship conference on

neural computation. It draws

preeminent academic

researchers from around the

Online Library Learning From Data Yaser Mostafa

world and is widely considered to be a showcase conference for new developments in network algorithms and architectures. The broad range of interdisciplinary research areas represented includes neural networks and genetic algorithms, cognitive science, neuroscience and biology, computer science, AI, applied mathematics, physics, and many branches of engineering. Only about 30% of the papers submitted are accepted for presentation at NIPS, so the quality is exceptionally high. All of the papers presented appear in these

Online Library Learning From Data Yaser Mostafa

proceedings.

Learning from Data
A Short Course
Machine Learning
The Art and Science of
Algorithms that Make Sense
of Data
Cambridge University
Press

This book describes how neural networks operate from the mathematical point of view. As a result, neural networks can be interpreted both as function universal approximators and information processors. The book bridges the gap between ideas and concepts of neural networks, which are used nowadays at an intuitive level, and the precise modern mathematical language, presenting the

Online Library Learning From Data Yaser Mostafa

best practices of the former and enjoying the robustness and elegance of the latter. This book can be used in a graduate course in deep learning, with the first few parts being accessible to senior undergraduates. In addition, the book will be of wide interest to machine learning researchers who are interested in a theoretical understanding of the subject.

Now in its second edition, D.S. Malik brings his proven approach to C++ programming to the CS2 course. Clearly written with the student in mind, this text focuses on Data Structures and includes advanced topics in C++ such

Online Library Learning From Data Yaser Mostafa

as Linked Lists and the Standard Template Library (STL). The text features abundant visual diagrams, examples, and extended Programming Examples, all of which serve to illuminate difficult concepts. Complete programming code and clear display of syntax, explanation, and example are used throughout the text, and each chapter concludes with a robust exercise set. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.