
Practical Artificial Intelligence For Dummies

Artificial Intelligence for HR
Practical Applications
Practical Machine Learning with H2O
Enterprise Artificial Intelligence Transformation
Machine Learning in Action
Algorithms For Dummies
A Plain English Introduction (Third Edition)
Deep Learning for Coders with fastai and PyTorch
Practical Artificial Intelligence with Swift
A Problem-Solver's Guide to Building Real-World Intelligent Systems
Artificial Intelligence For Dummies
Practical Applications
The AI Advantage
A Visual, Interactive Guide to Artificial Intelligence
Introducing Machine Learning
A non-technical, hype-free guide to prospering in the AI era
Practical Deep Learning for Cloud, Mobile, and Edge
The Book of Why
Zero to AI
Machine Learning For Dummies
An Enterprise Playbook
Artificial Intelligence in Finance
Practical Machine Learning with Python
Deep Learning Illustrated
Artificial Intelligence for Marketing
Artificial Intelligence and Deep Learning in Pathology
Artificial Intelligence in Basic
A Playbook for Business Value and Digital Transformation
Powerful, Scalable Techniques for Deep Learning and AI
Practical Machine Learning for Computer Vision
Practical Machine Learning for Data Analysis Using Python
Practical Artificial Intelligence
Practical Artificial Intelligence and Blockchain
Deep Learning For Dummies
Calculus of Thought
Neuromorphic Logistic Regression in Cognitive Machines
Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow
Artificial Intelligence in Practice

JORDON SUTTON

Artificial Intelligence for HR Cambridge University Press

Machine Learning For Dummies John Wiley & Sons

Practical Applications Simon and Schuster

Guides professionals and students through the rapidly growing field of machine learning with hands-on examples in the popular R programming language Machine learning—a branch of Artificial Intelligence (AI) which enables computers to improve their results and learn new approaches without explicit instructions—allows organizations to reveal patterns in their data and incorporate predictive analytics into their decision-making process. Practical Machine Learning in R provides a hands-on approach to solving business problems with intelligent, self-learning computer algorithms. Bestselling author and data analytics experts Fred Nwanganga and Mike Chapple explain what machine learning is, demonstrate its organizational benefits, and provide hands-on examples created in the R programming language. A perfect guide for professional self-taught learners or students in an introductory machine learning course, this reader-friendly book illustrates the numerous real-world business uses of machine learning approaches. Clear and detailed chapters cover data wrangling, R programming with the popular RStudio tool, classification and regression techniques, performance evaluation, and more. Explores data management techniques, including data collection, exploration and dimensionality reduction Covers unsupervised learning, where readers identify and summarize patterns using approaches such as apriori, eclat and clustering Describes the principles behind the Nearest Neighbor, Decision Tree and Naive Bayes classification techniques Explains how to evaluate and choose the right model, as well as how to improve model performance using ensemble methods such as Random Forest and XGBoost Practical Machine Learning in R is a must-have guide for business analysts, data scientists, and other professionals interested in leveraging the power of AI to solve business problems, as well as students and independent learners seeking to enter the field.

Practical Machine Learning with H2O Basic Books

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with

an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Enterprise Artificial Intelligence Transformation "O'Reilly Media, Inc."

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Machine Learning in Action Springer

Calculus of Thought: Neuromorphic Logistic Regression in Cognitive Machines is a must-read for all scientists about a very simple computation method designed to simulate big-data neural processing. This book is inspired by the Calculus Ratiocinator idea of Gottfried Leibniz, which is that machine computation should be developed to simulate human cognitive processes, thus avoiding problematic subjective bias in analytic solutions to practical and scientific problems. The reduced error logistic regression (RELR) method is proposed as such a "Calculus of Thought." This book reviews how RELR's completely automated processing may parallel important aspects of explicit and implicit

learning in neural processes. It emphasizes the fact that RELR is really just a simple adjustment to already widely used logistic regression, along with RELR's new applications that go well beyond standard logistic regression in prediction and explanation. Readers will learn how RELR solves some of the most basic problems in today's big and small data related to high dimensionality, multicollinearity, and cognitive bias in capricious outcomes commonly involving human behavior. Provides a high-level introduction and detailed reviews of the neural, statistical and machine learning knowledge base as a foundation for a new era of smarter machines Argues that smarter machine learning to handle both explanation and prediction without cognitive bias must have a foundation in cognitive neuroscience and must embody similar explicit and implicit learning principles that occur in the brain

Algorithms For Dummies John Wiley & Sons

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

A Plain English Introduction (Third Edition) "O'Reilly Media, Inc."

Summary How can artificial intelligence transform your business? In *Zero to AI*, you'll explore a variety of practical AI applications you can use to improve customer experiences, optimize marketing, help you cut costs, and more. In this engaging guide written for business leaders and technology pros alike, authors and AI experts Nicolò Valigi and Gianluca Mauro use fascinating projects, hands-on activities, and real-world explanations to make it clear how your business can benefit from AI. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology There's no doubt that artificial intelligence has made some impressive headlines recently, from besting chess and Go grand masters to producing uncanny deep fakes that blur the lines of reality. But what can AI do for you? If you want to understand how AI will impact your business before you invest your time and money, this book is for you. About the book *Zero to AI* uses clear examples and jargon-free explanations to show the practical benefits of AI. Each chapter explores a real-world case study demonstrating how

companies like Google and Netflix use AI to shape their industries. You begin at the beginning, with a primer on core AI concepts and realistic business outcomes. To help you prepare for the transition, the book breaks down a successful AI implementation, including advice on hiring the right team and making decisions about resources, risks, and costs. What's inside Identifying where AI can help your organization Designing an AI strategy Evaluating project scope and business impact Using AI to boost conversion rates, curate content, and analyze feedback Understanding how modern AI works and what it can/can't do About the reader For anyone who wants to gain an understanding of practical artificial intelligence and learn how to design and develop projects with high business impact. About the author Gianluca Mauro and Nicolò Valigi are the cofounders of AI Academy, a company specializing in AI trainings and consulting. Table of Contents: 1. An introduction to artificial intelligence PART 1 - UNDERSTANDING AI 2. Artificial intelligence for core business data 3. AI for sales and marketing 4. AI for media 5. AI for natural language 6. AI for content curation and community building PART 2 - BUILDING AI 7. Ready—finding AI opportunities 8. Set—preparing data, technology, and people 9. Go—AI implementation strategy 10. What lies ahead

Deep Learning for Coders with fastai and PyTorch Kogan Page Publishers

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Practical Artificial Intelligence with Swift John Wiley & Sons

Artificial Intelligence in BASIC presents some of the central ideas and practical applications of artificial intelligence (AI) using the BASIC programs. This eight-chapter book aims to explain these ideas of AI that can be used to produce programs on microcomputers. After providing an overview of the concept of AI, this book goes on examining the features and difficulties of a heuristic solution in a wide range of human problems. The discussion then shifts to the application of a heuristic solution to a two-ply search program for a two-person game. The following chapters are devoted to the other components of AI, including the expert systems, memory structure, pattern recognition, and language. The concluding chapter deals with the alternative and auxiliary approaches to the study of AI and its practical applications. Computer scientists and programmers will find this work invaluable. *A Problem-Solver's Guide to Building Real-World Intelligent Systems* John Wiley & Sons

This practical book shows you how to employ machine learning models to extract information from images. ML engineers and data scientists will learn how to solve a variety of image problems

including classification, object detection, autoencoders, image generation, counting, and captioning with proven ML techniques. This book provides a great introduction to end-to-end deep learning: dataset creation, data preprocessing, model design, model training, evaluation, deployment, and interpretability. Google engineers Valliappa Lakshmanan, Martin Görner, and Ryan Gillard show you how to develop accurate and explainable computer vision ML models and put them into large-scale production using robust ML architecture in a flexible and maintainable way. You'll learn how to design, train, evaluate, and predict with models written in TensorFlow or Keras. You'll learn how to: Design ML architecture for computer vision tasks Select a model (such as ResNet, SqueezeNet, or EfficientNet) appropriate to your task Create an end-to-end ML pipeline to train, evaluate, deploy, and explain your model Preprocess images for data augmentation and to support learnability Incorporate explainability and responsible AI best practices Deploy image models as web services or on edge devices Monitor and manage ML models

MIT Press

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Artificial Intelligence For Dummies "O'Reilly Media, Inc."

The widespread adoption of AI and machine learning is revolutionizing many industries today. Once these technologies are combined with the programmatic availability of historical and real-time financial data, the financial industry will also change fundamentally. With this practical book, you'll learn how to use AI and machine learning to discover statistical inefficiencies in financial markets and exploit them through algorithmic trading. Author Yves Hilpisch shows practitioners, students, and academics in both finance and data science practical ways to apply machine learning and deep learning algorithms to finance. Thanks to lots of self-contained Python examples, you'll be able to replicate all results and figures presented in the book. In five parts, this guide helps you: Learn central notions and algorithms from AI, including recent breakthroughs on the way to artificial general intelligence (AGI) and superintelligence (SI) Understand why data-driven finance, AI, and machine learning will have a lasting impact on financial theory and practice Apply neural networks and reinforcement learning to discover statistical inefficiencies in financial markets Identify and exploit economic inefficiencies through backtesting and algorithmic trading--the automated execution of trading strategies Understand how AI will influence the competitive dynamics in the financial industry and what the potential emergence of a financial singularity might bring about

Practical Applications Packt Publishing Ltd

Artificial Intelligence in Accounting: Practical Applications was written with a simple goal: to provide accountants with a foundational understanding of AI and its many business and accounting applications. It is meant to serve as a guide for identifying opportunities to implement AI initiatives to increase productivity and profitability. This book will help you answer questions about what AI is and how it is used in the accounting profession today. Offering practical guidance that you can leverage for your organization, this book provides an overview of essential AI concepts and technologies that accountants should know, such as machine learning, deep learning, and natural language processing. It also describes accounting-specific applications of robotic process automation and text mining. Illustrated with case studies and interviews with representatives from global professional services firms, this concise volume makes a significant contribution to examining the intersection of AI and the accounting profession. This innovative book also explores the challenges and ethical considerations of AI. It will be of great interest to accounting practitioners, researchers, educators, and students.

The AI Advantage Addison-Wesley Professional

A straightforward, non-technical guide to the next major marketing tool Artificial Intelligence for Marketing presents a tightly-focused introduction to machine learning, written specifically for marketing professionals. This book will not teach you to be a data scientist—but it does explain how Artificial Intelligence and Machine Learning will revolutionize your company's marketing strategy, and teach you how to use it most effectively. Data and analytics have become table stakes in modern marketing, but the field is ever-evolving with data scientists continually developing new algorithms—where does that leave you? How can marketers use the latest data science developments to their advantage? This book walks you through the "need-to-know" aspects of Artificial Intelligence, including natural language processing, speech recognition, and the power of Machine Learning to show you how to make the most of this technology in a practical, tactical way. Simple illustrations clarify complex concepts, and case studies show how real-world companies are taking the next leap forward. Straightforward, pragmatic, and with no math required, this book will help you: Speak intelligently about Artificial Intelligence and its advantages in marketing Understand how marketers without a Data Science degree can make use of machine learning technology Collaborate with data scientists as a subject matter expert to help develop focused-use applications Help your company gain a competitive advantage by leveraging leading-edge technology in marketing Marketing and data science are two fast-moving, turbulent spheres that often intersect; that intersection is where marketing professionals pick up the tools and methods to move their company forward. Artificial Intelligence and Machine Learning provide a data-driven basis for more robust and intensely-targeted marketing strategies—and companies that effectively utilize these latest tools will reap the benefit in the marketplace. Artificial Intelligence for Marketing provides a nontechnical crash course to help you stay ahead of the curve.

A Visual, Interactive Guide to Artificial Intelligence Packt Publishing Ltd

This accessible and engaging textbook presents a concise introduction to the exciting field of artificial intelligence (AI). The broad-ranging discussion covers the key subdisciplines within the field, describing practical algorithms and concrete applications in the areas of agents, logic, search,

reasoning under uncertainty, machine learning, neural networks, and reinforcement learning. Fully revised and updated, this much-anticipated second edition also includes new material on deep learning. Topics and features: presents an application-focused and hands-on approach to learning, with supplementary teaching resources provided at an associated website; contains numerous study exercises and solutions, highlighted examples, definitions, theorems, and illustrative cartoons; includes chapters on predicate logic, PROLOG, heuristic search, probabilistic reasoning, machine learning and data mining, neural networks and reinforcement learning; reports on developments in deep learning, including applications of neural networks to generate creative content such as text, music and art (NEW); examines performance evaluation of clustering algorithms, and presents two practical examples explaining Bayes' theorem and its relevance in everyday life (NEW); discusses search algorithms, analyzing the cycle check, explaining route planning for car navigation systems, and introducing Monte Carlo Tree Search (NEW); includes a section in the introduction on AI and society, discussing the implications of AI on topics such as employment and transportation (NEW). Ideal for foundation courses or modules on AI, this easy-to-read textbook offers an excellent overview of the field for students of computer science and other technical disciplines, requiring no more than a high-school level of knowledge of mathematics to understand the material.

Introducing Machine Learning Apress

Create and implement AI-based features in your Swift apps for iOS, macOS, tvOS, and watchOS. With this practical book, programmers and developers of all kinds will find a one-stop shop for AI and machine learning with Swift. Taking a task-based approach, you'll learn how to build features that use powerful AI features to identify images, make predictions, generate content, recommend things, and more. AI is increasingly essential for every developer—and you don't need to be a data scientist or mathematician to take advantage of it in your apps. Explore Swift-based AI and ML techniques for building applications. Learn where and how AI-driven features make sense. Inspect tools such as Apple's Python-powered Turi Create and Google's Swift for TensorFlow to train and build models. I: Fundamentals and Tools—Learn AI basics, our task-based approach, and discover how to build or find a dataset. II: Task Based AI—Build vision, audio, text, motion, and augmentation-related features; learn how to convert preexisting models. III: Beyond—Discover the theory behind task-based practice, explore AI and ML methods, and learn how you can build it all from scratch... if you want to

A non-technical, hype-free guide to prospering in the AI era John Wiley & Sons

Take a deep dive into deep learning Deep learning provides the means for discerning patterns in the data that drive online business and social media outlets. Deep Learning for Dummies gives you the information you need to take the mystery out of the topic—and all of the underlying technologies associated with it. In no time, you'll make sense of those increasingly confusing algorithms, and find a simple and safe environment to experiment with deep learning. The book develops a sense of precisely what deep learning can do at a high level and then provides examples of the major deep learning application types. Includes sample code Provides real-world examples within the approachable text Offers hands-on activities to make learning easier Shows you how to use Deep Learning more effectively with the right tools This book is perfect for those who want to better understand the basis of the underlying technologies that we use each and every day.

Practical Deep Learning for Cloud, Mobile, and Edge CRC Press

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of *Wait But Why* Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beylveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

The Book of Why "O'Reilly Media, Inc."

Discover how algorithms shape and impact our digital world All data, big or small, starts with algorithms. Algorithms are mathematical equations that determine what we see—based on our likes, dislikes, queries, views, interests, relationships, and more—online. They are, in a sense, the electronic gatekeepers to our digital, as well as our physical, world. This book demystifies the subject of algorithms so you can understand how important they are business and scientific decision making. Algorithms for Dummies is a clear and concise primer for everyday people who are interested in algorithms and how they impact our digital lives. Based on the fact that we already live in a world where algorithms are behind most of the technology we use, this book offers eye-opening information on the pervasiveness and importance of this mathematical science—how it plays out in our everyday digestion of news and entertainment, as well as in its influence on our social interactions and consumerism. Readers even learn how to program an algorithm using Python! Become well-versed in the major areas comprising algorithms Examine the incredible history behind algorithms Get familiar with real-world applications of problem-solving procedures Experience hands-on development of an algorithm from start to finish with Python If you have a nagging curiosity about why an ad for that hammock you checked out on Amazon is appearing on your

Facebook page, you'll find *Algorithm for Dummies* to be an enlightening introduction to this integral realm of math, science, and business.

Zero to AI "O'Reilly Media, Inc."

Recent advances in computational algorithms, along with the advent of whole slide imaging as a platform for embedding artificial intelligence (AI), are transforming pattern recognition and image interpretation for diagnosis and prognosis. Yet most pathologists have just a passing knowledge of data mining, machine learning, and AI, and little exposure to the vast potential of these powerful new tools for medicine in general and pathology in particular. In *Artificial Intelligence and Deep Learning in Pathology*, Dr. Stanley Cohen covers the nuts and bolts of all aspects of machine

learning, up to and including AI, bringing familiarity and understanding to pathologists at all levels of experience. Focuses heavily on applications in medicine, especially pathology, making unfamiliar material accessible and avoiding complex mathematics whenever possible. Covers digital pathology as a platform for primary diagnosis and augmentation via deep learning, whole slide imaging for 2D and 3D analysis, and general principles of image analysis and deep learning. Discusses and explains recent accomplishments such as algorithms used to diagnose skin cancer from photographs, AI-based platforms developed to identify lesions of the retina, using computer vision to interpret electrocardiograms, identifying mitoses in cancer using learning algorithms vs. signal processing algorithms, and many more.