

Kubeflow makes much easier. With this practical guide, data scientists, data engineers, and platform architects will learn how to plan and execute a Kubeflow project that can support workflows from on-premises to the cloud. Kubeflow is an open source Kubernetes-native platform based on Google's internal machine learning pipelines, and yet major cloud vendors including AWS and Azure advocate the use of Kubernetes and Kubeflow to manage containers and machine learning infrastructure. In today's cloud-based world, this book is ideal for any team planning to build machine learning applications. With this book, you will:

- Get a concise overview of Kubernetes and Kubeflow
- Learn how to plan and build a Kubeflow installation
- Operate, monitor, and automate your installation
- Provide your Kubeflow installation with adequate security
- Serve machine learning models on Kubeflow

Harness the power of the easy-to-use, high-performance fastai framework to rapidly create complete deep learning solutions with few lines of code

Key Features

- Discover how to apply state-of-the-art deep learning techniques to real-world problems
- Build and train neural networks using the power and flexibility of the fastai framework
- Use deep learning to tackle problems such as image classification and text classification

Book Description

fastai is an easy-to-use deep learning framework built on top of PyTorch that lets you rapidly create complete deep learning solutions with as few as 10 lines of code. Both predominant low-level deep learning frameworks, TensorFlow and PyTorch, require a lot of code, even for straightforward applications. In contrast, fastai handles the messy details for you and lets you focus on applying deep learning to actually solve problems. The book begins by summarizing the value of fastai and showing you how to create a simple 'hello world' deep learning application with fastai. You'll then learn how to use fastai for all four application areas that the framework explicitly supports: tabular data, text data (NLP), recommender systems, and vision data. As you advance, you'll work through a series of practical examples that illustrate how to create real-world applications of each type. Next, you'll learn how to deploy fastai models, including creating a simple web application that predicts what object is depicted in an image. The book wraps up with an overview of the advanced features of fastai. By the end of this fastai book, you'll be able to create your own deep learning applications using fastai. You'll also have learned how to use fastai to prepare raw datasets, explore datasets, train deep learning models, and deploy trained models. What you will learn

- Prepare real-world raw datasets to train fastai deep learning models
- Train fastai deep learning models using text and tabular data
- Create recommender systems with fastai
- Find out how to assess whether fastai is a good fit for a given problem
- Deploy fastai deep learning models in web applications
- Train fastai deep learning models for image classification

Who this book is for

This book is for data scientists, machine learning developers, and deep learning enthusiasts looking to explore the fastai framework using a recipe-based approach. Working knowledge of the Python programming language and machine learning basics is strongly recommended to get the most out of this deep learning

book.

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 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.

This book introduces a paradigm of reverse hypothesis machines (RHM), focusing on knowledge innovation and machine learning. Knowledge- acquisition -based learning is constrained by large volumes of data and is time consuming. Hence Knowledge innovation based learning is the need of time. Since under-learning results in cognitive inabilities and over-learning compromises freedom, there is need for optimal machine learning. All existing learning techniques rely on mapping input and output and establishing mathematical relationships between them. Though methods change the paradigm remains the same—the forward hypothesis machine paradigm, which tries to minimize uncertainty. The RHM, on the other hand, makes use of uncertainty for creative learning. The approach uses limited data to help identify new and surprising solutions. It focuses on improving learnability, unlike traditional approaches, which focus on accuracy. The book is useful as a reference book for machine learning researchers and professionals as well as machine intelligence enthusiasts. It can also used by practitioners to develop new machine learning applications to solve problems that require

creativity.

Artificial neural networks (ANNs) present many benefits in analyzing complex data in a proficient manner. As an effective and efficient problem-solving method, ANNs are incredibly useful in many different fields. From education to medicine and banking to engineering, artificial neural networks are a growing phenomenon as more realize the plethora of uses and benefits they provide. Due to their complexity, it is vital for researchers to understand ANN capabilities in various fields. The Research Anthology on Artificial Neural Network Applications covers critical topics related to artificial neural networks and their multitude of applications in a number of diverse areas including medicine, finance, operations research, business, social media, security, and more. Covering everything from the applications and uses of artificial neural networks to deep learning and non-linear problems, this book is ideal for computer scientists, IT specialists, data scientists, technologists, business owners, engineers, government agencies, researchers, academicians, and students, as well as anyone who is interested in learning more about how artificial neural networks can be used across a wide range of fields.

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

Information modeling and knowledge bases have become an important area of academic and industry research in the 21st century, addressing complexities of modeling that reach beyond the traditional borders of information systems and academic computer science research. This book presents 32 reviewed, selected and updated papers delivered at the 29th International Conference on Information Modeling and Knowledge Bases (EJC2019), held in Lappeenranta, Finland, from 3 to 7 June 2019. In addition, two papers based on the keynote presentations and one paper edited from the discussion of the panel session are included in the book. The conference provided a forum to exchange scientific results and experience, and attracted academics and practitioners working with information and knowledge. The papers cover a

wide range of topics, ranging from knowledge discovery through conceptual and linguistic modeling, knowledge and information modeling and discovery, cross-cultural communication and social computing, environmental modeling and engineering, and multimedia data modeling and systems to complex scientific problem-solving. The conference presentation sessions: Learning and Linguistics; Systems and Processes; Data and Knowledge Representation; Models and Interface; Formalizations and Reasoning; Models and Modeling; Machine Learning; Models and Programming; Environment and Predictions; and Emotion Modeling and Social Networks reflect the main themes of the conference. The book also includes 2 extended publications of keynote addresses: 'Philosophical Foundations of Conceptual Modeling' and 'Sustainable Solid Waste Management using Life Cycle Modeling for Environmental Impact Assessment', as well as additional material covering the discussion and findings of the panel session. Providing an overview of current research in the field, the book will be of interest to all those working with information systems, information modeling and knowledge bases.

Building models is a small part of the story when it comes to deploying machine learning applications. The entire process involves developing, orchestrating, deploying, and running scalable and portable machine learning workloads--a process Kubeflow makes much easier. This practical book shows data scientists, data engineers, and platform architects how to plan and execute a Kubeflow project to make their Kubernetes workflows portable and scalable. Authors Josh Patterson, Michael Katzenellenbogen, and Austin Harris demonstrate how this open source platform orchestrates workflows by managing machine learning pipelines. You'll learn how to plan and execute a Kubeflow platform that can support workflows from on-premises to cloud providers including Google, Amazon, and Microsoft. Dive into Kubeflow architecture and learn best practices for using the platform

- Understand the process of planning your Kubeflow deployment
- Install Kubeflow on an existing on-premises Kubernetes cluster
- Deploy Kubeflow on Google Cloud Platform step-by-step from the command line
- Use the managed Amazon Elastic Kubernetes Service (EKS) to deploy Kubeflow on AWS
- Deploy and manage Kubeflow across a network of Azure cloud data centers around the world
- Use KFServing to develop and deploy machine learning models

This book proposes complex hierarchical deep architectures (HDA) for predicting bankruptcy, a topical issue for business and corporate institutions that in the past has been tackled using statistical, market-based and machine-intelligence prediction models. The HDA are formed through fuzzy rough tensor deep staking networks (FRTDSN) with structured, hierarchical rough Bayesian (HRB) models. FRTDSN is formalized through TDSN and fuzzy rough sets, and HRB is formed by incorporating probabilistic rough sets in structured hierarchical Bayesian model. Then FRTDSN is integrated with HRB to form the compound FRTDSN-HRB model. HRB enhances the prediction accuracy of FRTDSN-HRB model.

The experimental datasets are adopted from Korean construction companies and American and European non-financial companies, and the research presented focuses on the impact of choice of cut-off points, sampling procedures and business cycle on the accuracy of bankruptcy prediction models. The book also highlights the fact that misclassification can result in erroneous predictions leading to prohibitive costs to investors and the economy, and shows that choice of cut-off point and sampling procedures affect rankings of various models. It also suggests that empirical cut-off points estimated from training samples result in the lowest misclassification costs for all the models. The book confirms that FRTDSN-HRB achieves superior performance compared to other statistical and soft-computing models. The experimental results are given in terms of several important statistical parameters revolving different business cycles and sub-cycles for the datasets considered and are of immense benefit to researchers working in this area.

As the 4th Industrial Revolution is restructuring human societal organization into, so-called, Society 5.0, the field of Machine Learning (and its sub-field of Deep Learning) and related technologies is growing continuously and rapidly, developing in both itself and towards applications in many other disciplines. Researchers worldwide aim at incorporating cognitive abilities into machines, such as learning and problem solving. When machines and software systems have been enhanced with Machine Learning/Deep Learning components, they become better and more efficient at performing specific tasks. Consequently, Machine Learning/Deep Learning stands out as a research discipline due to its worldwide pace of growth in both theoretical advances and areas of application, while achieving very high rates of success and promising major impact in science, technology and society. The book at hand aims at exposing its readers to some of the most significant Advances in Machine Learning/Deep Learning-based Technologies. The book consists of an editorial note and an additional ten (10) chapters, all invited from authors who work on the corresponding chapter theme and are recognized for their significant research contributions. In more detail, the chapters in the book are organized into five parts, namely (i) Machine Learning/Deep Learning in Socializing and Entertainment, (ii) Machine Learning/Deep Learning in Education, (iii) Machine Learning/Deep Learning in Security, (iv) Machine Learning/Deep Learning in Time Series Forecasting, and (v) Machine Learning in Video Coding and Information Extraction. This research book is directed towards professors, researchers, scientists, engineers and students in Machine Learning/Deep Learning-related disciplines. It is also directed towards readers who come from other disciplines and are interested in becoming versed in some of the most recent Machine Learning/Deep Learning-based technologies. An extensive list of bibliographic references at the end of each chapter guides the readers to probe further into the application areas of interest to them. How can machine learning--especially deep neural networks--make a real difference in your organization? This hands-on guide not only provides practical information, but helps you get started building efficient deep learning networks. The

authors provide the fundamentals of deep learning--tuning, parallelization, vectorization, and building pipelines--that are valid for any library before introducing the open source Deeplearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J.

Artificial intelligence and machine learning are considered as hot technologies of this century. As these technologies move from research labs to enterprise data centers, the need for skilled professionals is continuously on the rise. This book is intended for IT and business professionals looking to gain proficiency in these technologies but are turned off by the complex mathematical equations. This book is also useful for students in the area of artificial intelligence and machine learning to gain a conceptual understanding of the algorithms and get an industry perspective. This book is an ideal place to start your journey as

- Core concepts of machine learning algorithms are explained in plain English using illustrations, data tables and examples
- Intuitive meaning of the mathematics behind popular machine learning algorithms explained
- Covers classical machine learning, neural networks and deep learning algorithms

At a time when the IT industry is focusing on reskilling its vast human resources, Machine intelligence is a very timely publication. It has a simple approach that builds up from basics, which would help software engineers and students looking to learn about the field as well as those who might have started off without the benefit of a structured introduction or sound basics.

Highly recommended. - Siddhartha S, Founder and CEO of Intain - Financial technology startup
Suresh has written a very accessible book for practitioners. The book has depth yet avoids excessive mathematics. The coverage of the subject is very good and has most of the concepts required for understanding machine learning if someone is looking for depth. For senior management, it will provide a good overview. It is well written. I highly recommend it. - Whee Teck ONG, CEO of Trusted Source and VP of Singapore Computer Society

????????????????????,????????,????????,????,????,??,??,?????,???,?????,?????,????,???,???,?????????,?????????????????????
?

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun.

Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your

deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside

Training deep neural networks
Implementing modules and loss functions
Utilizing pretrained models from PyTorch Hub
Exploring code samples in Jupyter Notebooks

About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. consultant based in Munich, Germany and a PyTorch core developer.

Table of Contents

PART 1 - CORE PYTORCH

1 Introducing deep learning and the PyTorch Library

2 Pretrained networks

3 It starts with a tensor

4 Real-world data representation using tensors

5 The mechanics of learning

6 Using a neural network to fit the data

7 Telling birds from airplanes: Learning from images

8 Using convolutions to generalize

PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER

9 Using PyTorch to fight cancer

10 Combining data

sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Unlock the groundbreaking advances of deep learning with this extensively revised new edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In *Deep Learning with Python, Second Edition* you will learn: Deep learning from first principles Image classification and image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation *Deep Learning with Python* has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks.

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like automated language translation, image recognition, and more. Deep learning is quickly becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach—even if you have no background in mathematics or data science. This book shows you how to get started. About the book *Deep Learning with Python, Second Edition* introduces the field of deep learning using Python and the powerful Keras library. In this revised and expanded new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll quickly pick up the skills you need to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions

????????????,????????????:????????,????????,????????,????????????????????.

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 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.

?????:????

At the dawn of the 4th Industrial Revolution, the field of Deep Learning (a sub-field of Artificial Intelligence and Machine Learning) is growing continuously and rapidly, developing both theoretically and towards applications in increasingly many and diverse other disciplines. The book at hand aims at exposing its reader to some of the most significant recent advances in deep learning-based technological applications and consists of an editorial note and an additional fifteen (15) chapters. All chapters in the book were invited from authors who work in the corresponding chapter theme and are recognized for their significant research contributions. In more detail, the chapters in the book are organized into six parts, namely (1) Deep Learning in Sensing, (2) Deep Learning in Social Media and IOT, (3) Deep Learning in the Medical Field, (4) Deep Learning in Systems Control, (5) Deep Learning in Feature Vector Processing, and (6) Evaluation of Algorithm Performance. This research book is directed towards professors, researchers, scientists, engineers and students in computer science-related disciplines. It is also directed towards readers who come from other disciplines and are interested in becoming versed in some of the most recent deep learning-based technological applications. An extensive list of bibliographic references at the end of each chapter guides the readers to probe deeper into their application areas of interest.

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn

Google Translate, fluent conversations with Siri and Alexa, and enormous profits from automated trading on the New York Stock Exchange. Deep learning networks can play poker better than professional poker players and defeat a world champion at Go. In this book, Terry Sejnowski explains how deep learning went from being an arcane academic field to a disruptive technology in the information economy. Sejnowski played an important role in the founding of deep learning, as one of a small group of researchers in the 1980s who challenged the prevailing logic-and-symbol based version of AI. The new version of AI Sejnowski and others developed, which became deep learning, is fueled instead by data. Deep networks learn from data in the same way that babies experience the world, starting with fresh eyes and gradually acquiring the skills needed to navigate novel environments. Learning algorithms extract information from raw data; information can be used to create knowledge; knowledge underlies understanding; understanding leads to wisdom. Someday a driverless car will know the road better than you do and drive with more skill; a deep learning network will diagnose your illness; a personal cognitive assistant will augment your puny human brain. It took nature many millions of years to evolve human intelligence; AI is on a trajectory measured in decades. Sejnowski prepares us for a deep learning future.

This book shows machine learning enthusiasts and practitioners how to get the best of both worlds by deriving Fisher kernels from deep learning models. In addition, the book shares insight on how to store and retrieve large-dimensional Fisher vectors using feature selection and compression techniques. Feature selection and feature compression are two of the most popular off-the-shelf methods for reducing data's high-dimensional memory footprint and thus making it suitable for large-scale visual retrieval and classification. Kernel methods long remained the de facto standard for solving large-scale object classification tasks using low-level features, until the revival of deep models in 2006. Later, they made a comeback with improved Fisher vectors in 2010. However, their supremacy was always challenged by various versions of deep models, now considered to be the state of the art for solving various machine learning and computer vision tasks. Although the two research paradigms differ significantly, the excellent performance of Fisher kernels on the Image Net large-scale object classification dataset has caught the attention of numerous kernel practitioners, and many have drawn parallels between the two frameworks for improving the empirical performance on benchmark classification tasks. Exploring concrete examples on different data sets, the book compares the computational and statistical aspects of different dimensionality reduction approaches and identifies metrics to show which approach is superior to the other for Fisher vector encodings. It also provides references to some of the most useful resources that could provide practitioners and machine learning enthusiasts a quick start for learning and implementing a variety of deep learning models and kernel functions.

The present book is primarily intended for undergraduate and postgraduate students of computer science and engineering, information technology, and electrical and electronics engineering. It bridges the gaps in knowledge of the seemingly difficult areas of machine learning and nature inspired computing. The text is written in a highly interactive manner, which satisfies the learning curiosity of any reader. Content of the text has been diligently organized to offer seamless learning experience. The text begins

with introduction to machine learning, which is followed by explanation of different aspects of machine learning. Various supervised, unsupervised, reinforced and nature inspired learning techniques are included in the text book with numerous examples and case studies. Different aspects of new machine learning and nature inspired learning algorithms are explained in-depth. The well-explained algorithms and pseudo codes for each topic make this book useful for students. The book also throws light on areas like prediction and classification systems. Key Features • Day to day examples and pictorial representations for deeper understanding of the subject • Helps readers easily create programs/applications • Research oriented approach • More case studies and worked-out examples for each machine learning algorithm than any other book

Information in today's advancing world is rapidly expanding and becoming widely available. This eruption of data has made handling it a daunting and time-consuming task. Natural language processing (NLP) is a method that applies linguistics and algorithms to large amounts of this data to make it more valuable. NLP improves the interaction between humans and computers, yet there remains a lack of research that focuses on the practical implementations of this trending approach. Neural Networks for Natural Language Processing is a collection of innovative research on the methods and applications of linguistic information processing and its computational properties. This publication will support readers with performing sentence classification and language generation using neural networks, apply deep learning models to solve machine translation and conversation problems, and apply deep structured semantic models on information retrieval and natural language applications. While highlighting topics including deep learning, query entity recognition, and information retrieval, this book is ideally designed for research and development professionals, IT specialists, industrialists, technology developers, data analysts, data scientists, academics, researchers, and students seeking current research on the fundamental concepts and techniques of natural language processing. Volume 1 discusses various machine learning & cognitive science approaches, presenting high-throughput research by experts in this area. Bringing together machine learning, cognitive science and other aspects of artificial intelligence to help provide a roadmap for future research on intelligent systems, the book is a valuable reference resource for students, researchers and industry practitioners wanting to keep abreast of recent developments in this dynamic, exciting and profitable research field. Volume 2 provides a systematic and comprehensive overview of machine learning with cognitive science methods and technologies which have played an important role at the core of practical solutions for a wide scope of tasks between handheld apps, industrial process control, autonomous vehicles, environmental policies, life sciences, playing computer games, computational theory, and engineering development. The chapters in this book focus on readers interested in machine learning, cognitive and neuro-inspired computational systems theories, mechanisms, and architecture, which underline human and animal behaviour, and their application to conscious and intelligent systems. In the current version, it focuses on the successful implementation and step-by-step explanation of practical applications of the domain. It also offers a wide range of inspiring and interesting cutting-edge contributions to applications of machine learning and cognitive science such as healthcare products, medical electronics, and gaming. Overall, these two volumes provide valuable information on effective, cutting-edge techniques

and approaches for students, researchers, practitioners, and academicians working in the field of AI, neural network, machine learning, and cognitive science. Furthermore, the purpose of this book is to address the interests of a broad spectrum of practitioners, students, and researchers, who are interested in applying machine learning and cognitive science methods in their respective domains.

If you're an executive, manager, or anyone interested in leveraging AI within your organization, this is your guide. You'll understand exactly what AI is, learn how to identify AI opportunities, and develop and execute a successful AI vision and strategy. Alex Castrounis, business consultant and former IndyCar engineer and race strategist, examines the value of AI and shows you how to develop an AI vision and strategy that benefits both people and business. AI is exciting, powerful, and game changing—but too many AI initiatives end in failure. With this book, you'll explore the risks, considerations, trade-offs, and constraints for pursuing an AI initiative. You'll learn how to create better human experiences and greater business success through winning AI solutions and human-centered products. Use the book's AIPB Framework to conduct end-to-end, goal-driven innovation and value creation with AI Define a goal-aligned AI vision and strategy for stakeholders, including businesses, customers, and users Leverage AI successfully by focusing on concepts such as scientific innovation and AI readiness and maturity Understand the importance of executive leadership for pursuing AI initiatives "A must read for business executives and managers interested in learning about AI and unlocking its benefits. Alex Castrounis has simplified complex topics so that anyone can begin to leverage AI within their organization." - Dan Park, GM & Director, Uber "Alex Castrounis has been at the forefront of helping organizations understand the promise of AI and leverage its benefits, while avoiding the many pitfalls that can derail success. In this essential book, he shares his expertise with the rest of us." - Dean Wampler, Ph.D., VP, Fast Data Engineering at Lightbend

This book is a collection of the most recent approaches that combine metaheuristics and machine learning. Some of the methods considered in this book are evolutionary, swarm, machine learning, and deep learning. The chapters were classified based on the content; then, the sections are thematic. Different applications and implementations are included; in this sense, the book provides theory and practical content with novel machine learning and metaheuristic algorithms. The chapters were compiled using a scientific perspective. Accordingly, the book is primarily intended for undergraduate and postgraduate students of Science, Engineering, and Computational Mathematics and is useful in courses on Artificial Intelligence, Advanced Machine Learning, among others. Likewise, the book is useful for research from the evolutionary computation, artificial intelligence, and image processing communities.

This volume includes selected contributions presented during the 2nd edition of the international conference on WaterEnergyNEXUS which was held in Salerno, Italy in November 2018. This conference was organized by the Sanitary Environmental Engineering Division (SEED) of the University of Salerno (Italy) in cooperation with Advanced Institute of Water Industry at Kyungpook National University (Korea) and with The Energy and Resources Institute, TERI (India). The initiative received the patronage of UNESCO – World Water Association Programme (WWAP) and of the International Water Association (IWA) and was organized with the support of Springer (MENA Publishing Program), Arab Water Council (AWC), Korean Society of Environmental Engineering (KSEE) and Italian Society of Sanitary Environmental Engineering Professors (GITISA). With the support of international experts invited as plenary and keynote speakers, the conference aimed to give a

platform for Euro-Mediterranean countries to share and discuss key topics on such water-energy issues through the presentation of nature-based solutions, advanced technologies and best practices for a more sustainable environment. This volume gives a general and brief overview on current research focusing on emerging Water-Energy-Nexus issues and challenges and its potential applications to a variety of environmental problems that are impacting the Euro-Mediterranean zone and surrounding regions. A selection of novel and alternative solutions applied worldwide are included. The volume contains over about one hundred carefully refereed contributions from 44 countries worldwide selected for the conference. Topics covered include (1) Nexus framework and governance, (2) Environmental solutions for the sustainable development of the water sector, (3) future clean energy technologies and systems under water constraints, (4) environmental engineering and management, (5) Implementation and best practices Intended for researchers in environmental engineering, environmental science, chemistry, and civil engineering. This volume is also an invaluable guide for industry professionals working in both water and energy sectors.

This book provides IT professionals, educators, researchers, and students a compendium of knowledge on smart sensors and devices, types of sensors, data analysis and monitoring with the help of smart sensors, decision making, impact of machine learning algorithms, and artificial intelligence-related methodologies for data analysis and understanding of smart applications in networks. Smart sensor networks play an important role in the establishment of network devices which can easily interact with physical world through plethora of variety of sensors for collecting and monitoring the surrounding context and allowing environment information. Apart from military applications, smart sensor networks are used in many civilian applications nowadays and there is a need to manage high volume of demands in related applications. This book comprises of 9 chapters and presents a valuable insight on the original research and review articles on the latest achievements that contributes to the field of smart sensor networks and their usage in real-life applications like smart city, smart home, e-healthcare, smart social sensing networks, etc. Chapters illustrate technological advances and trends, examine research opportunities, highlight best practices and standards, and discuss applications and adoption. Some chapters also provide holistic and multiple perspectives while examining the impact of smart sensor networks and the role of data analytics, data sharing, and its control along with future prospects.

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural

Acces PDF Deep Learning A Practitioners Approach

networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

Deep LearningDL4J and BeyondO'Reilly Media

[Copyright: 0533e11ff8914b9a8a1412f14be28832](#)