

Fundamentals Of Neural Networks: Architectures, Algorithms And Applications: United States Edition (Pie)

As book review editor of the IEEE Transactions on Neural Networks, Mohamad Hassoun has had the opportunity to assess the multitude of books on artificial neural networks that have appeared in recent years. Now, in Fundamentals of Artificial Neural Networks, he provides the first systematic account of artificial neural network paradigms by identifying clearly the fundamental concepts and major methodologies underlying most of the current theory and practice employed by neural network researchers. Such a systematic and unified treatment, although sadly lacking in most recent texts on neural networks, makes the subject more accessible to students and practitioners. Here, important results are integrated in order to more fully explain a wide range of existing empirical observations and commonly used heuristics. There are numerous illustrative examples, over 200 end-of-chapter analytical and computer-based problems that will aid in the development of neural network analysis and design skills, and a bibliography of nearly 700 references. Proceeding in a clear and logical fashion, the first two chapters present the basic building blocks and concepts of artificial neural networks and analyze the computational capabilities of the basic network architectures involved. Supervised, reinforcement, and unsupervised learning rules in simple nets are brought together in a common framework in chapter three. The convergence and solution properties of these learning rules are then treated mathematically in chapter four, using the "average learning equation" analysis approach. This organization of material makes it natural to switch into learning multilayer nets using backprop and its variants, described in chapter five. Chapter six covers most of the major neural network paradigms, while associative memories and energy minimizing nets are given detailed coverage in the next chapter. The final chapter takes up Boltzmann machines and Boltzmann learning along with other global search/optimization algorithms such as stochastic gradient search, simulated annealing, and genetic algorithms.

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 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 Learning in Practice helps you learn how to develop and optimize a model for your projects using Deep Learning (DL) methods and architectures. Key features: Demonstrates a quick review on Python, NumPy, and TensorFlow fundamentals. Explains and provides examples of deploying TensorFlow and Keras in several projects. Explains the fundamentals of Artificial Neural Networks (ANNs). Presents several examples and applications of ANNs.

Learning the most popular DL algorithms features. Explains and provides examples for the DL algorithms that are presented in this book. Analyzes the DL network's parameter and hyperparameters. Reviews state-of-the-art DL examples. Necessary and main steps for DL modeling. Implements a Virtual Assistant Robot (VAR) using DL methods. Necessary and fundamental information to choose a proper DL algorithm. Gives instructions to learn how to optimize your DL model IN PRACTICE. This book is useful for undergraduate and graduate students, as well as practitioners in industry and academia. It will serve as a useful reference for learning deep learning fundamentals and implementing a deep learning model for any project, step by step.

New technologies in engineering, physics and biomedicine are demanding increasingly complex methods of digital signal processing. By presenting the latest research work the authors demonstrate how real-time recurrent neural networks (RNNs) can be implemented to expand the range of traditional signal processing techniques and to help combat the problem of prediction. Within this text neural networks are considered as massively interconnected nonlinear adaptive filters. Analyses the relationships between RNNs and various nonlinear models and filters, and introduces spatio-temporal architectures together with the concepts of modularity and nesting Examines stability and relaxation within RNNsPresents on-line learning algorithms for nonlinear adaptive filters and introduces new paradigms which exploit the concepts of a priori and a posteriori errors, data-reusing adaptation, and normalisation Studies convergence and stability of on-line learning algorithms based upon optimisation techniques such as contraction mapping and fixed point iteration Describes strategies for the exploitation of inherent relationships between parameters in RNNs Discusses practical issues such as predictability and nonlinearity detecting and includes several practical applications in areas such as air pollutant modelling and prediction, attractor discovery and chaos, ECG signal processing, and speech processing Recurrent Neural Networks for Prediction offers a new insight into the learning algorithms, architectures and stability of recurrent neural networks and, consequently, will have instant appeal. It provides an extensive background for researchers, academics and postgraduates enabling them to apply such networks in new applications. VISIT OUR COMMUNICATIONS TECHNOLOGY WEBSITE!

<http://www.wiley.co.uk/commstech/> VISIT OUR WEB PAGE! <http://www.wiley.co.uk/>

Learning Algorithms, Architectures and Stability

Designing Next-Generation Machine Intelligence Algorithms

A Textbook

Architectures, Algorithms, and Applications

Deep Learning with Python

Artificial Neural Networks with TensorFlow 2

Explains current co-design and co-optimization methodologies for building hardware neural networks and algorithms for machine learning applications This book focuses on how to build energy-efficient hardware for neural networks with learning capabilities—and provides co-design and co-optimization methodologies for building hardware neural networks that can learn. Presenting a complete picture from high-level algorithm to low-level implementation details, Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design also covers many fundamentals and essentials in neural networks (e.g., deep learning), as well as hardware implementation of neural networks. The book begins with an overview of neural networks. It then discusses algorithms for utilizing and training rate-based artificial neural networks. Next comes an introduction to various options for executing neural networks, ranging from general-purpose processors to specialized hardware, from digital accelerator to analog accelerator. A design example on building energy-efficient accelerator for adaptive dynamic programming with neural networks is also presented. An examination of fundamental concepts and popular learning algorithms for spiking neural networks follows that, along with a look at the hardware for spiking neural networks. Then comes a chapter offering readers three design examples (two of which are based on conventional CMOS, and one on emerging nanotechnology) to implement the learning algorithm found in the previous chapter. The book concludes with an outlook on the future of neural network hardware. Includes cross-layer survey of hardware accelerators for neuromorphic algorithms Covers the co-design of architecture and algorithms with emerging devices for much-improved computing efficiency Focuses on the co-design of algorithms and hardware, which is especially critical for using emerging devices, such as traditional memristors or diffusive

memristors, for neuromorphic computing Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design is an ideal resource for researchers, scientists, software engineers, and hardware engineers dealing with the ever-increasing requirement on power consumption and response time. It is also excellent for teaching and training undergraduate and graduate students about the latest generation neural networks with powerful learning capabilities.

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

Pattern recognizers evolve across the sections into perceptrons, a layer of perceptrons, multiple-layered perceptrons, functional link nets, and radial basis function networks. Other networks covered in the process are learning vector quantization networks, self-organizing maps, and recursive neural networks. Backpropagation is derived in complete detail for one and two hidden layers for both unipolar and bipolar sigmoid activation functions.

Your one-stop guide to learning and implementing artificial neural networks with Keras effectively Key Features Design and create neural network architectures on different domains using Keras Integrate neural network models in your applications using this highly practical guide Get ready for the future of neural networks through transfer learning and predicting multi network models Book Description Neural networks are used to solve a wide range of problems in different areas of AI and deep learning. Hands-On Neural Networks with Keras will start with teaching you about the core concepts of neural networks. You will delve into combining different neural network models and work with real-world use cases, including computer vision, natural language understanding, synthetic data generation, and many more. Moving on, you will become well versed with convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory (LSTM) networks, autoencoders, and generative adversarial networks (GANs) using real-world training datasets. We will examine how to use CNNs for image recognition, how to use reinforcement learning agents, and many more. We will dive into the specific architectures of various networks and then implement each of them in a hands-on manner using industry-grade frameworks. By the end of this book, you will be highly familiar with all prominent deep learning models and frameworks, and the options you have when applying deep learning to real-world scenarios and embedding artificial intelligence as the core fabric of your organization. What you will learn Understand the fundamental nature and workflow of predictive data modeling Explore how different types of visual and linguistic signals are processed by neural networks Dive into the mathematical and statistical ideas behind how networks learn from data Design and implement various neural networks such as CNNs, LSTMs, and GANs Use different architectures to tackle cognitive tasks and embed intelligence in systems Learn how to generate synthetic data and use augmentation strategies to improve your models Stay on top of the latest academic and commercial developments in the field of AI Who this book is for This book is for machine learning practitioners, deep learning researchers and AI enthusiasts who are looking to get well versed with different neural network architecture using Keras. Working knowledge of Python programming language is mandatory.

Design and create neural networks using deep learning and artificial intelligence principles

Making AI Less Susceptible to Adversarial Trickery

Volume 1

Deep Learning in Practice

A Practical Course

Foundations of Neural Networks

The past fifteen years has witnessed an explosive growth in the fundamental research and applications of artificial neural networks (ANNs) and fuzzy logic (FL). The main impetus behind this growth has been the ability of such methods to offer solutions not amenable to conventional techniques, particularly in application domains involving pattern recognition, prediction and control. Although the origins of ANNs and FL may be traced back to the 1940s and 1960s, respectively, the most rapid progress has only been achieved in the last fifteen years. This has been due to significant theoretical advances in our understanding of ANNs and FL, complemented by major technological developments in high-speed computing. In geophysics, ANNs and FL have enjoyed significant success and are now employed routinely in the following areas (amongst others): 1. Exploration Seismology. (a) Seismic data processing (trace editing; first break picking; deconvolution and multiple suppression; wavelet estimation; velocity analysis; noise identification/reduction; statics analysis; dataset matching/prediction, attenuation), (b) AVO analysis, (c) Chimneys, (d) Compression I dimensionality reduction, (e) Shear-wave analysis, (f) Interpretation (event tracking; lithology prediction and well-log analysis; prospect appraisal; hydrocarbon prediction; inversion; reservoir characterisation; quality assessment; tomography). 2. Earthquake Seismology and Subterranean Nuclear Explosions. 3. Mineral Exploration. 4. Electromagnetic I Potential Field Exploration. (a) Electromagnetic methods, (b) Potential field methods, (c) Ground penetrating radar, (d) Remote sensing, (e) inversion.

State of the Art in Neural Networks and Their Applications presents the latest advances in artificial neural networks and their applications across a wide range of clinical diagnoses. Advances in the role of machine learning, artificial intelligence, deep learning, cognitive image processing and suitable data analytics useful for clinical diagnosis and research applications are covered, including relevant case studies. The application of Neural Network, Artificial Intelligence, and Machine Learning methods in biomedical image analysis have resulted in the development of computer-aided diagnostic (CAD) systems that aim towards the automatic early detection of several severe diseases. State of the Art in Neural Networks and Their Applications is presented in two volumes. Volume 1 covers the state-of-the-art deep learning approaches for the detection of renal, retinal, breast, skin, and dental abnormalities and more. Includes applications of neural networks, AI, machine learning, and deep learning techniques to a variety of imaging technologies Provides in-depth technical coverage of computer-aided diagnosis (CAD), with coverage of computer-aided classification, Unified Deep Learning Frameworks, mammography, fundus imaging, optical coherence tomography, cryo-electron tomography, 3D MRI, CT, and more. Covers deep learning for several medical conditions including renal, retinal, breast, skin, and dental abnormalities, Medical Image Analysis, as well as detection, segmentation, and classification via AI.

As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn't trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this

adversarial input, and methods for increasing AI robustness to these attacks. If you're a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input Investigate methods used to generate adversarial input capable of fooling DNNs Explore real-world scenarios and model the adversarial threat Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data Examine some ways in which AI might become better at mimicking human perception in years to come

Provides an introduction to the neural network modeling of complex cognitive and neuropsychological processes. Over the past few years, computer modeling has become more prevalent in the clinical sciences as an alternative to traditional symbol-processing models. This book provides an introduction to the neural network modeling of complex cognitive and neuropsychological processes. It is intended to make the neural network approach accessible to practicing neuropsychologists, psychologists, neurologists, and psychiatrists. It will also be a useful resource for computer scientists, mathematicians, and interdisciplinary cognitive neuroscientists. The editors (in their introduction) and contributors explain the basic concepts behind modeling and avoid the use of high-level mathematics. The book is divided into four parts. Part I provides an extensive but basic overview of neural network modeling, including its history, present, and future trends. It also includes chapters on attention, memory, and primate studies. Part II discusses neural network models of behavioral states such as alcohol dependence, learned helplessness, depression, and waking and sleeping. Part III presents neural network models of neuropsychological tests such as the Wisconsin Card Sorting Task, the Tower of Hanoi, and the Stroop Test. Finally, part IV describes the application of neural network models to dementia: models of acetylcholine and memory, verbal fluency, Parkinsons disease, and Alzheimer's disease. Contributors J. Wesson Ashford, Rajendra D. Badgaiyan, Jean P. Banquet, Yves Burnod, Nelson Butters, John Cardoso, Agnes S. Chan, Jean-Pierre Changeux, Kerry L. Coburn, Jonathan D. Cohen, Laurent Cohen, Jose L. Contreras-Vidal, Antonio R. Damasio, Hanna Damasio, Stanislas Dehaene, Martha J. Farah, Joaquin M. Fuster, Philippe Gaussier, Angelika Gissler, Dylan G. Harwood, Michael E. Hasselmo, J. Allan Hobson, Sam Leven, Daniel S. Levine, Debra L. Long, Roderick K. Mahurin, Raymond L. Ownby, Randolph W. Parks, Michael I. Posner, David P. Salmon, David Servan-Schreiber, Chantal E. Stern, Jeffrey P. Sutton, Lynette J. Tippett, Daniel Tranel, Bradley Wyble

Neuropsychology and Cognitive Neuroscience

Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design

Recurrent Neural Networks for Prediction

From Fundamentals to Complex Pattern Recognition

Introduction to Deep Learning

Deep Learning

An introduction to neural networks written at an elementary level, with the new student in mind. The text features systematic discussions of the major neural networks and gives numerous examples, exercises and also 25 computer projects.

This book covers the details of concepts, the methods and the challenges of evolutionary deep neural networks design. The authors

begin by providing a brief introduction to deep neural networks, evolutionary computation. They also include some representative examples of both. Then they move on to describing the scope of evolutionary deep neural network design, and the fundamental methods of evolutionary deep neural network architecture design. Finally, they highlight the main challenges and some potential research directions in this emerging topic.

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Computational intelligence is a well-established paradigm, where new theories with a sound biological understanding have been evolving. The current experimental systems have many of the characteristics of biological computers (brains in other words) and are beginning to be built to perform a variety of tasks that are difficult or impossible to do with conventional computers. As evident, the

ultimate achievement in this field would be to mimic or exceed human cognitive capabilities including reasoning, recognition, creativity, emotions, understanding, learning and so on. This book comprising of 17 chapters offers a step-by-step introduction (in a chronological order) to the various modern computational intelligence tools used in practical problem solving. Starting with different search techniques including informed and uninformed search, heuristic search, minmax, alpha-beta pruning methods, evolutionary algorithms and swarm intelligent techniques; the authors illustrate the design of knowledge-based systems and advanced expert systems, which incorporate uncertainty and fuzziness. Machine learning algorithms including decision trees and artificial neural networks are presented and finally the fundamentals of hybrid intelligent systems are also depicted. Academics, scientists as well as engineers engaged in research, development and application of computational intelligence techniques, machine learning and data mining would find the comprehensive coverage of this book invaluable.

Fundamentals of Neural Networks : Architectures Algorithms and Applications

Deep Learning Illustrated

Neural Network Design

Theory and Applications

An Introduction to Neural Information Retrieval

Pattern Recognition Using Neural Networks

"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 Beyleveld 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.

Providing detailed examples of simple applications, this new book introduces the use of neural networks. It covers simple neural nets for pattern classification; pattern association; neural networks based on competition; adaptive-resonance theory; and more. For professionals working with neural networks.

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 Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its

operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

Neural Networks and Deep Learning

Fundamentals of Artificial Neural Networks

Theory and Algorithms for Engineers and Scientists

A Practitioner's Approach

A Modern Approach

ANN Architecture Machine Learning Projects

In this book, highly qualified multidisciplinary scientists grasp their recent researches motivated by the importance of artificial neural networks. It addresses advanced applications and innovative case studies for the next-generation optical networks based on modulation recognition using artificial neural networks, hardware ANN for gait generation of multi-legged robots, production of high-resolution soil property ANN maps, ANN and dynamic factor models to combine forecasts, ANN parameter recognition of engineering constants in Civil Engineering, ANN electricity consumption and generation forecasting, ANN for advanced process control, ANN breast cancer detection, ANN applications in biofuels, ANN modeling for manufacturing process optimization, spectral interference correction using a large-size spectrometer and ANN-based deep learning, solar radiation ANN prediction using NARX model, and ANN data assimilation for an atmospheric general circulation model.

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you
About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how

to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application. In response to the exponentially increasing need to analyze vast amounts of data, Neural Networks for Applied Sciences and Engineering: From Fundamentals to Complex Pattern Recognition provides scientists with a simple but systematic introduction to neural networks. Beginning with an introductory discussion on the role of neural networks in

This book provides comprehensive coverage of neural networks, their evolution, their structure, the problems they can solve, and their applications. The first half of the book looks at theoretical investigations on artificial neural networks and addresses the key architectures that are capable of implementation in various application scenarios. The second half is designed specifically for the production of solutions using artificial neural networks to solve practical problems arising from different areas of knowledge. It also describes the various implementation details that were taken into account to achieve the reported results. These aspects contribute to the maturation and improvement of experimental techniques to specify the neural network architecture that is most appropriate for a particular application scope. The book is appropriate for students in graduate and upper undergraduate courses in addition to researchers and professionals.

Artificial Intelligence with Python
Concepts, Methods, and Challenges
Intelligent Systems
Geophysical Applications of Artificial Neural Networks and Fuzzy Logic
Fundamentals of Deep Learning
Artificial Neural Networks

Forecasting is required in many situations. Stocking an inventory may require forecasts of demand months in advance. Telecommunication routing requires traffic forecasts a few minutes ahead. Whatever the circumstances or time horizons involved, forecasting is an important aid in effective and efficient planning. This textbook provides a comprehensive introduction to forecasting methods and presents enough information about each method for readers to use them sensibly.

Efficient Query Processing for Scalable Web Search will be a valuable reference for researchers and developers working on This tutorial provides an accessible, yet comprehensive, overview of the state-of-the-art of Neural Information Retrieval. Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their 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. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types

with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop Develop machine learning models across various domains. This book offers a single source that provides comprehensive coverage of the capabilities of TensorFlow 2 through the use of realistic, scenario-based projects. After learning what's new in TensorFlow 2, you'll dive right into developing machine learning models through applicable projects. This book covers a wide variety of ANN architectures—starting from working with a simple sequential network to advanced CNN, RNN, LSTM, DCGAN, and so on. A full chapter is devoted to each kind of network and each chapter consists of a full project describing the network architecture used, the theory behind that architecture, what data set is used, the pre-processing of data, model training, testing and performance optimizations, and analysis. This practical approach can either be used from the beginning through to the end or, if you're already familiar with basic ML models, you can dive right into the application that interests you. Line-by-line explanations on major code segments help to fill in the details as you work and the entire project source is available to you online for learning and further experimentation. With Artificial Neural Networks with TensorFlow 2 you'll see just how wide the range of TensorFlow's capabilities are. What You'll Learn Develop Machine Learning Applications Translate languages using neural networks Compose images with style transfer Who This Book Is For Beginners, practitioners, and hard-cored developers who want to master machine and deep learning with TensorFlow 2. The reader should have working concepts of ML basics and terminologies.

Multivariate Statistical Machine Learning Methods for Genomic Prediction

Hands-On Neural Networks with Keras

Neural Networks in the Analysis and Design of Structures

From Logical Calculus to Artificial Intelligence

Object-Oriented Neural Networks in C++

State of the Art in Neural Networks and Their Applications

Neural Networks are a new, interdisciplinary tool for information processing. Neurocomputing being successfully introduced to structural problems which are difficult or even impossible to be analysed by standard computers (hard computing). The book is devoted to foundations and

applications of NNs in the structural mechanics and design of structures.

This book introduces readers to the fundamentals of deep neural network architectures, with a special emphasis on memristor circuits and systems. At first, the book offers an overview of neuro-memristive systems, including memristor devices, models, and theory, as well as an introduction to deep learning neural networks such as multi-layer networks, convolution neural networks, hierarchical temporal memory, and long short term memories, and deep neuro-fuzzy networks. It then focuses on the design of these neural networks using memristor crossbar architectures in detail. The book integrates the theory with various applications of neuro-memristive circuits and systems. It provides an introductory tutorial on a range of issues in the design, evaluation techniques, and implementations of different deep neural network architectures with memristors.

New technologies in engineering, physics and biomedicine are demanding increasingly complex methods of digital signal processing. By presenting the latest research work the authors demonstrate how real-time recurrent neural networks (RNNs) can be implemented to expand the range of traditional signal processing techniques and to help combat the problem of prediction. Within this text neural networks are considered as massively interconnected nonlinear adaptive filters. Analyses the relationships between RNNs and various nonlinear models and filters, and introduces spatio-temporal architecture.

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights

that may spark new ideas.

Evolutionary Deep Neural Network Design

Strengthening Deep Neural Networks

Neural Networks for Applied Sciences and Engineering

Fundamentals of Neural Networks: Architectures, Algorithms and Applications

A Visual, Interactive Guide to Artificial Intelligence

Advanced Applications for Artificial Neural Networks

This book is open access under a CC BY 4.0 license This open access book brings together the latest genome base prediction models currently being used by statisticians, breeders and data scientists. It provides an accessible way to understand the theory behind each statistical learning tool, the required pre-processing, the basics of model building, how to train statistical learning methods, the basic R scripts needed to implement each statistical learning tool, and the output of each tool. To do so, for each tool the book provides background theory, some elements of the R statistical software for its implementation, the conceptual underpinnings, and at least two illustrative examples with data from real-world genomic selection experiments. Lastly, worked-out examples help readers check their own comprehension. The book will greatly appeal to readers in plant (and animal) breeding, geneticists and statisticians, as it provides in a very accessible way the necessary theory, the appropriate R code, and illustrative examples for a complete understanding of each statistical learning tool. In addition, it weighs the advantages and disadvantages of each tool.

This book provides a broad yet detailed introduction to neural networks and machine learning in a statistical framework. A single, comprehensive resource for study and further research, it explores the major popular neural network models and statistical learning approaches with examples and exercises and allows readers to gain a practical working understanding of the content. This updated new edition presents recently published results and includes six new chapters that correspond to the recent advances in computational learning theory, sparse coding, deep learning, big data and cloud computing. Each chapter features state-of-the-art descriptions and significant research findings. The topics covered include: • multilayer perceptron; • the Hopfield network; • associative memory models; • clustering models and algorithms; • the radial basis function network; • recurrent neural networks; • nonnegative matrix factorization; • independent component analysis; • probabilistic and Bayesian networks; and •

fuzzy sets and logic. Focusing on the prominent accomplishments and their practical aspects, this book provides academic and technical staff, as well as graduate students and researchers with a solid foundation and comprehensive reference on the fields of neural networks, pattern recognition, signal processing, and machine learning.

"This book is distinctive in that it implements nodes and links as base objects and then composes them into four different kinds of neural networks. Roger's writing is clear....The text and code are both quite readable. Overall, this book will be useful to anyone who wants to implement neural networks in C++ (and, to a lesser extent, in other object-oriented programming languages)...I recommend this book to anyone who wants to implement neural networks in C++."--D.L. Chester, Newark, Delaware in COMPUTING REVIEWS

Object-Oriented Neural Networks in C++ is a valuable tool for anyone who wants to understand, implement, or utilize neural networks. This book/disk package provides the reader with a foundation from which any neural network architecture can be constructed. The author has employed object-oriented design and object-oriented programming concepts to develop a set of foundation neural network classes, and shows how these classes can be used to implement a variety of neural network architectures with a great deal of ease and flexibility. A wealth of neural network formulas (with standardized notation), object code implementations, and examples are provided to demonstrate the object-oriented approach to neural network architectures and to facilitate the development of new neural network architectures. This is the first book to take full advantage of the reusable nature of neural network classes.

Key Features

- * Describes how to use the classes provided to implement a variety of neural network architectures including ADALINE, Backpropagation, Self-Organizing, and BAM*
- * Provides a set of reusable neural network classes, created in C++, capable of implementing any neural network architecture*
- * Includes an IBM disk of the source code for the classes, which is platform independent*
- * Includes an IBM disk with C++ programs described in the book*

Neural Networks and Statistical Learning

Fundamentals of Neural Networks

Efficient Processing of Deep Neural Networks

Deep Learning Classifiers with Memristive Networks

Forecasting: principles and practice

An Introduction to Neural Networks