

Learning Deep Architectures For Ai

This book presents a wealth of deep-learning algorithms and demonstrates their design process. It also highlights the need for a prudent alignment with the essential characteristics of the nature of learning encountered in the practical problems being tackled. Intended for readers interested in acquiring practical knowledge of analysis, design, and deployment of deep learning solutions to real-world problems, it covers a wide range of the paradigm’s algorithms and their applications in diverse areas including imaging, seismic tomography, smart grids, surveillance and security, and health care, among others. Featuring systematic and comprehensive discussions on the development processes, their evaluation, and relevance, the book offers insights into fundamental design strategies for algorithms of deep learning.

This book offers a timely reflection on the remarkable range of algorithms and applications that have made the area of deep learning so attractive and heavily researched today. Introducing the diversity of learning mechanisms in the environment of big data, and presenting authoritative studies in fields such as sensor design, health care, autonomous driving, industrial control and wireless communication, it enables readers to gain a practical understanding of design. The book also discusses systematic design procedures, optimization techniques, and validation processes.

This book introduces readers to the fundamental concepts of deep learning and offers practical insights into how this learning paradigm supports automatic mechanisms of structural knowledge representation. It discusses a number of multilayer architectures giving rise to tangible and functionally meaningful pieces of knowledge, and shows how the structural developments have become essential to the successful delivery of competitive practical solutions to real-world problems. The book also demonstrates how the architectural developments, which arise in the setting of deep learning, support detailed learning and refinements to the system design. Featuring detailed descriptions of the current trends in the design and analysis of deep learning topologies, the book offers practical guidelines and presents competitive solutions to various areas of language modeling, graph representation, and forecasting.

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.

A Visual, Interactive Guide to Artificial Intelligence

4th International Conference of Pioneering Computer Scientists, Engineers and Educators, ICPCSEE 2018, Zhengzhou, China, September 21-23, 2018, Proceedings, Part II

31st Canadian Conference on Artificial Intelligence, Canadian AI 2018, Toronto, ON, Canada, May 8-11, 2018, Proceedings

Deep In-memory Architectures for Machine Learning

Select Proceedings of ICAAAIML 2020

Development and Analysis of Deep Learning Architectures

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.

This book constitutes the refereed proceedings of the 29th Australasian Joint Conference on Artificial Intelligence, AI 2016, held in Hobart, TAS, Australia, in December 2016. The 40 full papers and 18 short papers presented together with 8 invited short papers were carefully reviewed and selected from 121 submissions. The papers are organized in topical sections on agents and multiagent systems; AI applications and innovations; big data; constraint satisfaction, search and optimisation; knowledge representation and reasoning; machine learning and data mining; social intelligence; and text mining and NLP. The proceedings also contains 2 contributions of the AI 2016 doctoral consortium and 6 contributions of the SMA 2016.

Theoretical results suggest that in order to learn the kind of complex functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

Learning Deep Architectures for AINow Publishers Inc

Multi-faceted Deep Learning

11th IFIP WG 10.3 International Conference, NPC 2014, Ilan, Taiwan, September 18-20, 2014, Proceedings

16th International Conference, ICAISC 2017, Zakopane, Poland, June 11-15, 2017, Proceedings, Part I

TensorFlow for Deep Learning

Multidisciplinary Approaches to Neural Computing

Dynamic Fuzzy Machine Learning

This book constitutes the refereed proceedings of the 31th Canadian Conference on Artificial Intelligence, Canadian AI 2018, held in Toronto, ON, Canada, in May 2018. The 16 regular papers and 18 short papers presented together with 7 Graduate Student Symposium papers and 4 Industry Track papers were carefully reviewed and selected from 72 submissions. The focus of the conference was on artificial intelligence research and advanced information and communications technology.

Concepts, tools, and techniques to explore deep learning architectures and methodologies
Key Features
Explore advanced deep learning architectures using various datasets and frameworks
Implement deep architectures for neural network models such as CNN, RNN, GAN, and many more
Discover design patterns and different challenges for various deep learning architectures
Book Description
Deep learning architectures are composed of multilevel nonlinear operations that represent high-level abstractions; this allows you to learn useful feature representations from the data. This book will help you learn and implement deep learning architectures to resolve various deep learning research problems. Hands-On Deep Learning Architectures with Python explains the essential learning algorithms used for deep and shallow architectures. Packed with practical implementations and ideas to help you build efficient artificial intelligence systems (AI), this book will help you learn how neural networks play a major role in building deep architectures. You will understand various deep learning architectures (such as AlexNet, VGG Net, GoogleNet) with easy-to-follow code and diagrams. In addition to this, the book will also guide you in building and training various deep architectures such as the Boltzmann mechanism, autoencoders, convolutional neural networks (CNNs), recurrent neural networks (RNNs), natural language processing (NLP), GAN, and more—all with practical implementations. By the end of this book, you will be able to construct deep models using popular frameworks and datasets with the required design patterns for each architecture. You will be ready to explore the potential of deep architectures in today's world. What you will learn
Implement CNNs, RNNs, and other commonly used architectures with Python
Explore architectures such as VGGNet, AlexNet, and GoogleNet
Build deep learning architectures for AI applications such as face and image recognition, fraud detection, and many more
Understand the architectures and applications of Boltzmann machines and autoencoders with concrete examples
Master artificial intelligence and neural network concepts and apply them to your architecture
Understand deep learning architectures for mobile and embedded systems
Who this book is for
If you’re a data scientist, machine learning developer/engineer, or deep learning practitioner, or are curious about AI and want to upgrade your knowledge of various deep learning architectures, this book will appeal to you. You are expected to have some knowledge of statistics and machine learning algorithms to get the best out of this book

The two-volume set LNAI 10245 and LNAI 10246 constitutes the refereed proceedings of the 16th International Conference on Artificial Intelligence and Soft Computing, ICAISC 2017, held in Zakopane, Poland in June 2017. The 133 revised full papers presented were carefully reviewed and selected from 274 submissions. The papers included in the first volume are organized in the following five parts: neural networks and their applications; fuzzy systems and their applications; evolutionary algorithms and their applications; computer vision, image and speech analysis; and bioinformatics, biometrics and medical applications.

"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.

Artificial Intelligence for Sustainable Development: Theory, Practice and Future Applications

Artificial Intelligence Applications and Innovations

Artificial Intelligence and Deep Learning in Pathology

Models and Data

From Linear Regression to Reinforcement Learning

Applications of Artificial Intelligence and Machine Learning

This book presents a collection of contributions in the field of Artificial Neural Networks (ANNs). The themes addressed are multidisciplinary in nature, and closely connected in their ultimate aim to identify features from dynamic realistic signal exchanges and invariant machine representations that can be exploited to improve the quality of life of their end users. Mathematical tools like ANNs are currently exploited in many scientific domains because of their solid theoretical background and effectiveness in providing solutions to many demanding tasks such as appropriately processing (both for extracting features and recognizing) mono- and bi-dimensional dynamic signals, solving strong nonlinearities in the data and providing general solutions for deep and fully connected architectures. Given the multidisciplinary nature of their use and the interdisciplinary characterization of the problems they are applied to - which range from medicine to psychology, industrial and social robotics, computer vision, and signal processing (among many others) - ANNs may provide a basis for redefining the concept of information processing. These reflections are supported by theoretical models and applications presented in the chapters of this book. This book is of primary importance for: (a) the academic research community, (b) the ICT market, (c) PhD students and early-stage researchers, (d) schools, hospitals, rehabilitation and assisted-living centers, and (e) representatives of multimedia industries and standardization bodies.

This book describes the recent innovation of deep in-memory architectures for realizing AI systems that operate at the edge of energy-latency-accuracy trade-offs. From first principles to lab prototypes, this book provides a comprehensive view of this emerging topic for both the practicing engineer in industry and the researcher in academia. The book is a journey into the exciting world of AI systems in hardware.

Deep learning systems have gotten really great at identifying patterns in text, images, and video. But applications that create realistic images, natural sentences and paragraphs, or native-quality translations have proven elusive. Generative Adversarial Networks, or GANs, offer a promising solution to these challenges by pairing two competing neural networks' one that generates content and the other that rejects samples that are of poor quality. GANs in Action: Deep learning with Generative Adversarial Networks teaches you how to build and train your own generative adversarial networks. First, you'll get an introduction to generative modelling and how GANs work, along with an overview of their potential uses. Then, you'll start building your own simple adversarial system, as you explore the foundation of GAN architecture: the generator and discriminator networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Provides an overview of general deep learning methodology and its applications to a variety of signal and information processing tasks

Efficient Processing of Deep Neural Networks

Deep Learning for Natural Language Processing

A Practitioner's Approach

AI 2016: Advances in Artificial Intelligence

Tools for Engagement

Machine Learning and Information Processing

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 constitutes the proceedings of the 11th IFIP WG 10.3 International Conference on Network and Parallel Computing, NPC 2014, held in Ilan, Taiwan, in September 2014. The 42 full papers and 24 poster papers presented were carefully reviewed and selected from 196 submissions. They are organized in topical sections on systems, networks, and architectures, parallel and multi-core technologies, virtualization and cloud computing technologies, applications of parallel and distributed computing, and I/O, file systems, and data management.

This handbook presents some of the most recent topics in neural information processing, covering both theoretical concepts and practical applications. The contributions include: Deep architectures
Recurrent, recursive, and graph neural networks
Cellular neural networks
Bayesian networks
Approximation capabilities of neural networks
Semi-supervised learning
Statistical relational learning
Kernel methods for structured data
Multiple classifier systems
Self organisation and modal learning
Applications to content-based image retrieval, text mining in large document collections, and bioinformatics
This book is thought particularly for graduate students, researchers and practitioners, willing to deepen their knowledge on more advanced connectionist models and related learning paradigms.

The leading experts in system change and learning, with their school-based partners around the world, have created this essential companion to their runaway best-seller, Deep Learning: Engage the World Change the World. This hands-on guide provides a roadmap for building capacity in teachers,

schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. Dive Into Deep Learning: Tools for Engagement is rich with resources educators need to construct and drive meaningful deep learning experiences in order to develop the kind of mindset and know-how that is crucial to becoming a problem-solving change agent in our global society. Designed in full color, this easy-to-use guide is loaded with tools, tips, protocols, and real-world examples. It includes:

- A framework for deep learning that provides a pathway to develop the six global competencies needed to flourish in a complex world – character, citizenship, collaboration, communication, creativity, and critical thinking.
- Learning progressions to help educators analyze student work and measure progress.
- Learning design rubrics, templates and examples for incorporating the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital.
- Conditions rubrics, teacher self-assessment tools, and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts. Learn about, improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create learning experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

Neural Networks: Tricks of the Trade

Network and Parallel Computing

Proceedings of the 12th Annual Meeting of the BICA Society

Artificial Intelligence and Soft Computing

Deep Learning: Algorithms and Applications

Deep Learning

Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

Recent advances in computational algorithms, along with the advent of whole slide imaging as a platform for embedding artificial intelligence (AI), are transforming pattern recognition and image interpretation for diagnosis and prognosis. Yet most pathologists have just a passing knowledge of data mining, machine learning, and AI, and little exposure to the vast potential of these powerful new tools for medicine in general and pathology in particular. In Artificial Intelligence and Deep Learning in Pathology, Dr. Stanley Cohen covers the nuts and bolts of all aspects of machine learning, up to and including AI, bringing familiarity and understanding to pathologists at all levels of experience. Focuses heavily on applications in medicine, especially pathology, making unfamiliar material accessible and avoiding complex mathematics whenever possible. Covers digital pathology as a platform for primary diagnosis and augmentation via deep learning, whole slide imaging for 2D and 3D analysis, and general principles of image analysis and deep learning. Discusses and explains recent accomplishments such as algorithms used to diagnose skin cancer from photographs, AI-based platforms developed to identify lesions of the retina, using computer vision to interpret electrocardiograms, identifying mitoses in cancer using learning algorithms vs. signal processing algorithms, and many more.

This book constitutes the refereed proceedings of the 28th Canadian Conference on Artificial Intelligence, Canadian AI 2015, held in Halifax, Nova Scotia, Canada, in June 2015. The 15 regular papers and 12 short papers presented together with 8 papers from the Graduate Student Symposium were carefully reviewed and selected from 81 submissions. The papers are organized in topical sections such as agents, uncertainty and games; AI applications; NLP, text and social media mining; data mining and machine learning.

The book focuses on original approaches intended to support the development of biologically inspired cognitive architectures. It bridges together different disciplines, including artificial intelligence, linguistics, neuro- and social sciences, psychology and philosophy of mind, among others. The chapters are based on contributions presented at the 12th Annual Meeting of the BICA Society (BICA 2021), which consisted of two parallel virtual events: Information in Biologically Inspired Cognitive Architectures based Systems, held during the 2021 Summit of the International Society for the Study of Information, on September 12-19, 2021, from Vienna, Austria, and the 2021 International Workshop on Biologically Inspired Cognitive Architectures, held during the 21st ACM International Conference on Intelligent Virtual Agents, on September 14-17, 2021, from the Fukuchiyama City, Kyoto, Japan. The book discusses emerging methods, theories and ideas towards the realization of general-purpose humanlike artificial intelligence or fostering a better understanding of the ways the human mind works. It provides engineers, mathematicians, psychologists, computer scientists and other experts with a timely snapshot of recent research and a source of inspiration for future developments in the broadly intended areas of artificial intelligence and biological inspiration.

Artificial Intelligence in the Age of Neural Networks and Brain Computing

Advances in Deep Learning

Create deep neural networks to solve computational problems using TensorFlow and Keras

Advances in Artificial Intelligence

Proceedings of ICMLIP 2020

Artificial Intelligence Applications in Specialty Crops

Machine learning is widely used for data analysis. Dynamic fuzzy data are one of the most difficult types of data to analyse in the field of big data, cloud computing, the Internet of Things, and quantum information. At present, the processing of this kind of data is not very mature. The authors carried out more than 20 years of research, and show in this book their most important results. The seven chapters of the book are devoted to key topics such as dynamic fuzzy machine learning models, dynamic fuzzy self-learning subspace algorithms, fuzzy decision tree learning, dynamic concepts based on dynamic fuzzy sets, semi-supervised multi-task learning based on dynamic fuzzy data, dynamic fuzzy hierarchy learning, examination of multi-agent learning model based on dynamic fuzzy logic. This book can be used as a reference book for senior college students and graduate students as well as college teachers and scientific and technical personnel involved in computer science, artificial intelligence, machine learning, automation, data analysis, mathematics, management, cognitive science, and finance. It can be also used as the basis for teaching the principles of dynamic fuzzy learning.

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

This book covers algorithmic and hardware implementation techniques to enable embedded deep learning. The authors describe synergetic design approaches on the application-, algorithmic-, computer architecture-, and circuit-level that will help in achieving the goal of reducing the computational cost of deep learning algorithms. The impact of these techniques is displayed in four silicon prototypes for embedded deep learning. Gives a wide overview of a series of effective solutions for energy-efficient neural networks on battery constrained wearable devices; Discusses the optimization of neural networks for embedded deployment on all levels of the design hierarchy – applications, algorithms, hardware architectures, and circuits – supported by real silicon prototypes; Elaborates on how to design efficient Convolutional Neural Network processors, exploiting parallelism and data-reuse, sparse operations, and low-precision computations; Supports the introduced theory and design concepts by four real silicon prototypes. The physical realization's implementation and achieved performances are discussed elaborately to illustrate and highlight the introduced cross-layer design concepts.

This book includes selected papers from the 2nd International Conference on Machine Learning and Information Processing (ICMLIP 2020), held at Vardhaman College of Engineering, Jawaharlal Nehru Technological University (JNTU), Hyderabad, India, from November 28 to 29, 2020. It presents the latest developments and technical solutions in the areas of advanced computing and data sciences, covering machine learning, artificial intelligence, human-computer interaction, IoT, deep learning, image processing and pattern recognition, and signal and speech processing.

Big Data Analytics and Artificial Intelligence Against COVID-19: Innovation Vision and Approach

29th Australasian Joint Conference, Hobart, TAS, Australia, December 5-8, 2016, Proceedings

Develop Deep Learning Models for your Natural Language Problems

Methods and Applications

Deep Learning and the Game of Go

Dive Into Deep Learning

Artificial intelligence (AI) in its various forms — machine learning, chatbots, robots, agents, etc. — is increasingly being seen as a core component of enterprise business workflow and information management systems. The current promise and hype around AI are being driven by software vendors, academic research projects, and startups. However, we posit that the greatest promise and potential of AI is in its applications touching all organizational facets. With increasing business process and workflow maturity, coupled with recent trends in cloud computing, datafication, IoT, cybersecurity, and advanced analytics, there is an understanding that the challenges of tomorrow cannot be solely addressed by today's people, processes, and products. There is still considerable mystery, hype, and fear about a considerable amount of current discourse focuses on a dystopian future that could adversely affect humanity. Such opinions, with understandable fear of the unknown, don't consider the history of human innovation, the current state of business and technology, or the primarily augmentative nature of tomorrow's AI. This book demystifies AI for the enterprise. It takes readers from the basics (of multi-industry journey, and concludes with expert advice on everything an organization must do to succeed. Along the way, we debunk myths, provide practical pointers, and include best practices with applicable vignettes. AI brings to enterprise the capabilities that promise new ways by which professionals can address both mundane and interesting challenges more efficiently, effectively, and cost-effectively. AI also provides an opportunity for tomorrow's enterprise is to augment existing teams and resources with the power of AI in order to gain competitive advantage, discover new business models, establish or optimize new revenues, and achieve better customer and user satisfaction.

This book constitutes the refereed proceedings of the 15th IFIP WG 12.5 International Conference on Artificial Intelligence Applications and Innovations, AIAI 2019, held in Hersonissos, Crete, Greece, in May 2019. The 49 full papers and 6 short papers presented were carefully reviewed and selected from 101 submissions. They cover a broad range of topics such as deep learning ANN: genetic algorithms; reinforcement learning; recommender systems; intelligent telecommunications modeling; and intelligent hybrid systems using Internet of Things. The papers are organized in the following topical sections: AI anomaly detection - active learning; autonomous vehicles - aerial vehicles; biomedical AI; classification - clustering; constraint programming - brain inspired modeling; deep learning - convolutional ANN: fuzzy automata - logic based reasoning; machine learning - natural language; multi agent - IoT; nature inspired flight and robot; control - machine vision; and recommendation systems.

Artificial Intelligence in the Age of Neural Networks and Brain Computing demonstrates that existing disruptive implications and applications of AI is a development of the unique attributes of neural networks, mainly machine learning, distributed architectures, massive parallel processing, black-box inference, intrinsic nonlinearity and smart autonomous search engines. The book covers the major building blocks behind AI, provides a framework to deep learning, and launches novel and intriguing paradigms as future alternatives. The success of AI-based commercial products proposed by top industry leaders, such as Google, IBM, Microsoft, Intel and Amazon can be interpreted using this book. Developed from the 30th anniversary of the International Neural Network Society (INNS) and the 2017 International Joint Conference on Artificial Intelligence (IJCAI) and the 2017 International Conference on Artificial Intelligence and Applications (ICAAI), this book is a valuable resource for researchers in the field of Artificial Intelligence (AI) and Neural Networks (NN). Authored by top experts, global field pioneers and researchers working on cutting-edge applications in signal processing, speech recognition, games, adaptive control and decision-making Edited by high-level academics and researchers in intelligent systems and neural networks

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 workflows. 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 architectures and applications. 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 framework for processing data natively on Spark and Hadoop

Neural Networks for Natural Language Processing

Deep Learning Illustrated

Handbook on Neural Information Processing

15th IFIP WG 12.5 International Conference, AIAI 2019, Hersonissos, Crete, Greece, May 24–26, 2019, Proceedings

Deep learning with Generative Adversarial Networks

GANs in Action

This two volume set (CCIS 901 and 902) constitutes the refereed proceedings of the 4th International Conference of Pioneering Computer Scientists, Engineers and Educators, ICPCSEE 2018 (originally ICYCSEE) held in Zhengzhou, China, in September 2018. The 125 revised full papers presented in these two volumes were carefully reviewed and selected from 1057 submissions. The papers cover a wide range of topics related to basic theory and techniques for data science including mathematical issues in data science, computational theory for data science, big data management and applications, data quality and data preparation, evaluation and measurement in data science, data visualization, big data mining and knowledge management, infrastructure for data science, machine learning for data science, data security and privacy, applications of data science, case study of data science, multimedia data management and analysis, data-driven scientific research, data-driven bioinformatics, data-driven healthcare, data-driven management, data-driven eGovernment, data-driven smart city/planet, data marketing and economics, social media and recommendation systems, data-driven security, data-driven business model innovation, social and/or organizational impacts of data science.

This book includes research articles and expository papers on the applications of artificial intelligence and big data analytics to battle the pandemic. In the context of COVID-19, this book focuses on how big data analytic and artificial intelligence help fight COVID-19. The book is divided into four parts. The first part discusses the forecasting and visualization of the COVID-19 data. The second part describes applications of artificial intelligence in the COVID-19 diagnosis of chest X-Ray imaging. The third part discusses the insights of artificial intelligence to stop spread of COVID-19, while the last part presents deep learning and big data analytics which help fight the COVID-19.

This book covers a large set of methods in the field of Artificial Intelligence - Deep Learning applied to real-world problems. The fundamentals of the Deep Learning approach and different types of Deep Neural Networks (DNNs) are first summarized in this book, which offers a comprehensive preamble for further problem-oriented chapters. The most interesting and open problems of machine learning in the framework of Deep Learning are discussed in this book and solutions are proposed. This book illustrates how to implement the zero-shot learning with Deep Neural Network Classifiers, which require a large amount of training data. The lack of annotated training data naturally pushes the researchers to implement low supervision algorithms. Metric learning is a long-term research but in the framework of Deep Learning approaches, it gets freshness and originality. Fine-grained classification with a low inter-class variability is a difficult problem for any classification tasks. This book presents how it is solved, by using different modalities and attention mechanisms in 3D convolutional networks. Researchers focused on Machine Learning, Deep learning, Multimedia and Computer Vision will want to buy this book. Advanced level students studying computer science within these topic areas will also find this book useful.

The book presents a collection of peer-reviewed articles from the International Conference on Advances and Applications of Artificial Intelligence and Machine Learning - ICAAAIML 2020. The book covers research in artificial intelligence, machine learning, and deep learning applications in healthcare, agriculture, business, and security. This volume contains research papers from academicians, researchers as well as students. There are also papers on core concepts of computer networks, intelligent system design and deployment, real-time systems, wireless sensor networks, sensors and sensor nodes, software engineering, and image processing. This book will be a valuable resource for students, academics, and practitioners in the industry working on AI applications.

A Playbook for Business Value and Digital Transformation

Hands-On Deep Learning Architectures with Python

28th Canadian Conference on Artificial Intelligence, Canadian AI 2015, Halifax, Nova Scotia, Canada, June 2-5, 2015, Proceedings

Algorithms, Architectures and Circuits for Always-on Neural Network Processing

Data Science

The twenty last years have been marked by an increase in available data and computing power. In parallel to this trend, the focus of neural network research and the practice of training neural networks has undergone a number of important changes, for example, use of deep learning machines. The second edition of the book augments the first edition with more tricks, which have resulted from 14 years of theory and experimentation by some of the world's most prominent neural network researchers. These tricks can make a substantial difference (in terms of speed, ease of implementation, and accuracy) when it comes to putting algorithms to work on real problems.

This book introduces readers to both basic and advanced concepts in deep network models. It covers state-of-the-art deep architectures that many researchers are currently using to overcome the limitations of the traditional artificial neural networks. Various deep architecture models and their components are discussed in detail, and subsequently illustrated by algorithms and selected applications. In addition, the book explains in detail the transfer learning approach for faster training of deep models; the approach is also demonstrated on large volumes of fingerprint and face image datasets. In closing, it discusses the unique set of problems and challenges associated with these models.

This book highlights the latest advances in the field of artificial intelligence and related technologies, with a special focus on sustainable development and environmentally friendly artificial intelligence applications. Discussing theory, applications and research, it covers all aspects of artificial intelligence in the context of sustainable development.

Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning

Deep Learning: Concepts and Architectures
Biologically Inspired Cognitive Architectures 2021
Demystifying AI for the Enterprise
Embedded Deep Learning
Learning Deep Architectures for AI