Algorithms And Architectures Of Artificial Intelligence Frontiers In Artificial Intelligence And Applications

Provides an overview of methods developed in artificial intelligence for search, learning, problem solving and decision making. This book also gives an overview of algorithms and architectures of artificial intelligence that have reached the degree of maturity when a method can be presented as an algorithm.

Presents extensions of papers reported in the proceedings of the 1989 SPIE symposium. In the applications group articles cover the use of expert systems in computer aided education, hydrocarbon (oil) exploration, satellite image analysis for oceanography, particle beam accelerator tuning, design of internal combustion engines. Papers in the theory group cover theories, algorithms, architectures and software tools that can be used for modules within future systems.

Biologically Inspired Networking and Sensing: Algorithms and Architectures offers current perspectives and trends in biologically inspired networking, exploring various approaches aimed at improving network paradigms. Research contained within this compendium of research papers and surveys introduces researches in the fields of communication networks, performance modeling, and distributed computing to new advances in networking.

Hybrid architecture for intelligent systems is a new field of artificial intelligence concerned with the development

of the next generation of intelligent systems. This volume is the first book to delineate current research interests in hybrid architectures for intelligent systems. The book is divided into two parts. The first part is devoted to the theory, methodologies, and algorithms of intelligent hybrid systems. The second part examines current applications of intelligent hybrid systems in areas such as data analysis, pattern classification and recognition, intelligent robot control, medical diagnosis, architecture, wastewater treatment, and flexible manufacturing systems. Hybrid Architectures for Intelligent Systems is an important reference for computer scientists and electrical engineers involved with artificial intelligence, neural networks, parallel processing, robotics, and systems architecture.

MATLAB enables the design of artificial intelligence models through three essential pillars: Machine Learning, Deep Learning and Data Science. Using MATLAB, engineers and other experts have deployed thousands of machine learning applications. Automated machine learning (AutoML) generate automatically functionalities from training data and optimize models using hyperparameter fitting techniques such as Bayesian optimization. Use specialized functionalities extraction techniques, such as wavelet dispersion for signal or image data, and functionalities selection techniques, such as neighbor component analysis (NCA) or sequential functionalities selection. Deep Learning is a subset of machine learning based on artificial neural networks. The process of this learning is called deep because this network structure consists of having

multiple inputs, outputs and hidden layers. Each layer contains units that transform the input data into information, and in this way, the next layer can use it for a certain predictive task. In this way, a machine can learn through its own data processing. MATLAB has the tool Neural Network Toolbox (Deep Leraning toolbox fron release 18) that provides algorithms, functions, and apps to create, train, visualize, and simulate neural networks. You can perform classification, regression, clustering, dimensionality reduction, time-series forecasting, and dynamic system modeling and control. This book develops the neural network architectures used in Deep Learning.

Master the approaches and principles of Artificial Intelligence (AI) algorithms, and apply them to Data Science projects with Python and Julia code. Aspiring and practicing Data Science and AI professionals, along with Python and Julia programmers, will practice numerous AI algorithms and develop a more holistic understanding of the field of AI, and will learn when to use each framework to tackle projects in our increasingly complex world. The first two chapters introduce the field, with Chapter 1 surveying Deep Learning models and Chapter 2 providing an overview of algorithms beyond Deep Learning, including Optimization, Fuzzy Logic, and Artificial Creativity. The next chapters focus on AI frameworks; they contain data and Python and Julia code in a provided Docker, so you can practice. Chapter 3 covers Apache's MXNet, Chapter 4 covers TensorFlow, and Chapter 5 investigates Keras. After covering these Deep Learning frameworks, we explore a

series of optimization frameworks, with Chapter 6 covering Particle Swarm Optimization (PSO), Chapter 7 on Genetic Algorithms (GAs), and Chapter 8 discussing Simulated Annealing (SA). Chapter 9 begins our exploration of advanced AI methods, by covering Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). Chapter 10 discusses optimization ensembles and how they can add value to the Data Science pipeline. Chapter 11 contains several alternative AI frameworks including Extreme Learning Machines (ELMs), Capsule Networks (CapsNets), and Fuzzy Inference Systems (FIS). Chapter 12 covers other considerations complementary to the AI topics covered, including Big Data concepts, Data Science specialization areas, and useful data resources to experiment on. A comprehensive glossary is included, as well as a series of appendices covering Transfer Learning, Reinforcement Learning, Autoencoder Systems, and Generative Adversarial Networks. There is also an appendix on the business aspects of AI in data science projects, and an appendix on how to use the Docker image to access the book's data and code. The field of AI is vast, and can be overwhelming for the newcomer to approach. This book will arm you with a solid understanding of the field, plus inspire you to explore further.

Computer vision is one of the most complex and computationally intensive problem. Like any other computationally intensive problems, parallel pro cessing has been suggested as an approach to solving the problems in com puter vision. Computer vision employs

algorithms from a wide range of areas such as image and signal processing, advanced mathematics, graph theory, databases and artificial intelligence. Hence, not only are the comput ing requirements for solving vision problems tremendous but they also demand computers that are efficient to solve problems exhibiting vastly dif ferent characteristics. With recent advances in VLSI design technology, Single Instruction Multiple Data (SIMD) massively parallel computers have been proposed and built. However, such architectures have been shown to be useful for solving a very limited subset of the problems in vision. Specifically, algorithms from low level vision that involve computations closely mimicking the architec ture and require simple control and computations are suitable for massively parallel SIMD computers. An Integrated Vision System (IVS) involves com putations from low to high level vision to be executed in a systematic fashion and repeatedly. The interaction between computations and information dependent nature of the computations suggests that architectural require ments for computer vision systems can not be satisfied by massively parallel SIMD computers.

"This book gives an overview of methods developed in artificial intelligence for search, learning, problem solving and decision-making. It gives an overview of algorithms and architectures of artificial intelligence that have reached the degree of maturity when a method can be presented as an algorithm, or when a well-defined architecture is known, e.g. in neural nets and intelligent agents. It can be used as a handbook for a wide audience of application developers who are interested in using artificial intelligence methods in their

software products. Parts of the text are rather independent, so that one can look into the index and go directly to a description of a method presented in the form of an abstract algorithm or an architectural solution. The book can be used also as a textbook for a course in applied artificial intelligence. Exercises on the subject are added at the end of each chapter. Neither programming skills nor specific knowledge in computer science are expected from the reader. However, some parts of the text will be fully understood by those who know the terminology of computing well."

The field of neural information processing has two main objects: investigation into the functioning of biological neural networks and use of artificial neural networks to sol ve real world problems. Even before the reincarnation of the field of artificial neural networks in mid nineteen eighties, researchers have attempted to explore the engineering of human brain function. After the reincarnation, we have seen an emergence of a large number of neural network models and their successful applications to solve real world problems. This volume presents a collection of recent research and developments in the field of neural information processing. The book is organized in three Parts, i.e., (1) architectures, (2) learning algorithms, and (3) applications. Artificial neural networks consist of simple processing elements called neurons, which are connected by weights. The number of neurons and how they are connected to each other defines the architecture of a particular neural network. Part 1 of the book has nine chapters, demonstrating some of recent neural network architectures derived either to mimic aspects of human brain function or applied in some real world problems. Muresan provides a simple neural network model, based on spiking neurons that make use of shunting inhibition, which is capable of resisting small scale changes of stimulus. Hoshino and Zheng simulate a neural network of the auditory cortex to

investigate neural basis for encoding and perception of vowel sounds.

Neural Network Analysis, Architectures and Applications discusses the main areas of neural networks, with each authoritative chapter covering the latest information from different perspectives. Divided into three parts, the book first lays the groundwork for understanding and simplifying networks. It then describes novel architectures and algorithms, including pulse-stream techniques, cellular neural networks, and multiversion neural computing. The book concludes by examining various neural network applications, such as neuron-fuzzy control systems and image compression. This final part of the book also provides a case study involving oil spill detection. This book is invaluable for students and practitioners who have a basic understanding of neural computing yet want to broaden and deepen their knowledge of the field.

This is the third in a series of conferences devoted primarily to the theory and applications of artificial neural networks and genetic algorithms. The first such event was held in Innsbruck, Austria, in April 1993, the second in Ales, France, in April 1995. We are pleased to host the 1997 event in the mediaeval city of Norwich, England, and to carryon the fine tradition set by its predecessors of providing a relaxed and stimulating environment for both established and emerging researchers working in these and other, related fields. This series of conferences is unique in recognising the relation between the two main themes of artificial neural networks and genetic algorithms, each having its origin in a natural process fundamental to life on earth, and each now well established as a paradigm fundamental to continuing technological development through the solution of complex, industrial, commercial and financial problems. This is well illustrated in this volume by the numerous applications of both paradigms

to new and challenging problems. The third key theme of the series, therefore, is the integration of both technologies, either through the use of the genetic algorithm to construct the most effective network architecture for the problem in hand, or, more recently, the use of neural networks as approximate fitness functions for a genetic algorithm searching for good solutions in an 'incomplete' solution space, i.e. one for which the fitness is not easily established for every possible solution instance.

This volume is the first diverse and comprehensive treatment of algorithms and architectures for the realization of neural network systems. It presents techniques and diverse methods in numerous areas of this broad subject. The book covers major neural network systems structures for achieving effective systems, and illustrates them with examples. This volume includes Radial Basis Function networks, the Expand-and-Truncate Learning algorithm for the synthesis of Three-Layer Threshold Networks, weight initialization, fast and efficient variants of Hamming and Hopfield neural networks, discrete time synchronous multilevel neural systems with reduced VLSI demands, probabilistic design techniques, time-based techniques, techniques for reducing physical realization requirements, and applications to finite constraint problems. A unique and comprehensive reference for a broad array of algorithms and architectures, this book will be of use to practitioners, researchers, and students in industrial, manufacturing, electrical, and mechanical engineering, as well as in computer science and engineering. Key Features * Radial Basis Function networks * The Expand-and-Truncate Learning algorithm for the synthesis of Three-Layer Threshold Networks * Weight initialization * Fast and efficient variants of Hamming and Hopfield neural networks * Discrete time synchronous multilevel neural systems with reduced VLSI demands * Probabilistic design techniques * Time-based

techniques * Techniques for reducing physical realization requirements * Applications to finite constraint problems * Practical realization methods for Hebbian type associative memory systems * Parallel self-organizing hierarchical neural network systems * Dynamics of networks of biological neurons for utilization in computational neuroscience Practitioners, researchers, and students in industrial, manufacturing, electrical, and mechanical engineering, as well as in computer science and engineering, will find this volume a unique and comprehensive reference to a broad array of algorithms and architectures

This book starts by presenting the basics of reinforcement learning using highly intuitive and easy-to-understand examples and applications, and then introduces the cutting-edge research advances that make reinforcement learning capable of out-performing most state-of-art systems, and even humans in a number of applications. The book not only equips readers with an understanding of multiple advanced and innovative algorithms, but also prepares them to implement systems such as those created by Google Deep Mind in actual code. This book is intended for readers who want to both understand and apply advanced concepts in a field that combines the best of two worlds – deep learning and reinforcement learning – to tap the potential of 'advanced artificial intelligence' for creating real-world applications and game-winning algorithms.

Algorithms and Architectures of Artificial IntelligenceIOS Press

The authors pointed out several issues that should be considered as part of an evaluation of potential AI machine architectures. They have then exemplified several activities of the advanced architecture community that focus on improving the run-time

performance of AI systems. One can conclude that when the architecture matches the algorithm, the performance increase is substantial over that of von Neumann architectures. For example, when ASPRO is given an exact match forward chaining production system to execute, the performance increase is linear over von Neumann. If that same machine architecture were presented many procedure calls during execution, performance would degenerate to that of a sequential architecture. Real-time AI system designers should address the issue of matching their algorithms to machine architectures from the beginning of the system development process. The higher the real-time performance requirements become, the more critical this ongoing evaluation becomes.

This book explores the subject of artificial psychology and how the field must adapt human neuro-psychological testing techniques to provide adequate cognitive testing of advanced artificial intelligence systems. It shows how classical testing methods will reveal nothing about the cognitive nature of the systems and whether they are learning, reasoning, and evolving correctly; for these systems, the authors outline how testing techniques similar to/adapted from human psychological testing must be adopted, particularly in understanding how the system reacts to failure or relearning something it has learned incorrectly or inferred incorrectly. The authors provide insights into future architectures/capabilities that artificial cognitive systems will possess and how we can evaluate how well they are functioning. It discusses at length the notion of human/AI communication and

collaboration and explores such topics as knowledge development, knowledge modeling and ambiguity management, artificial cognition and self-evolution of learning, artificial brain components and cognitive architecture, and artificial psychological modeling. Explores the concepts of Artificial Psychology and Artificial Neuroscience as applied to advanced artificially cognitive systems; Provides insight into the world of cognitive architectures and biologically-based computing designs which will mimic human brain functionality in artificial intelligent systems of the future; Provides description and design of artificial psychological modeling to provide insight into how advanced artificial intelligent systems are learning and evolving; Explores artificial reasoning and inference architectures and the types of modeling and testing that will be required to "trust" an autonomous artificial intelligent systems.

This two volume set LNCS 8630 and 8631 constitutes the proceedings of the 14th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2014, held in Dalian, China, in August 2014. The 70 revised papers presented in the two volumes were selected from 285 submissions. The first volume comprises selected papers of the main conference and papers of the 1st International Workshop on Emerging Topics in Wireless and Mobile Computing, ETWMC 2014, the 5th International Workshop on Intelligent Communication Networks, IntelNet 2014, and the 5th International Workshop on Wireless Networks and Multimedia, WNM 2014. The second volume comprises selected papers of the main conference and papers of

the Workshop on Computing, Communication and Control Technologies in Intelligent Transportation System, 3C in ITS 2014, and the Workshop on Security and Privacy in Computer and Network Systems, SPCNS 2014.

The IEEE Third International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP-97) will be held in Melbourne, Australia from December 8th to 12th, 1997. The purpose of this important conference is to bring together developers and researchers from universities, industry and government to advance science and technology in distributed and parallel systems and processing.

Neural Networks for Optimization and Signal Processing A. Cichocki Warsaw University of Technology Poland R. Unbehauen Universität Erlangen-Nürnberg Germany Artificial neural networks can be employed to solve a wide spectrum of problems in optimization, parallel computing, matrix algebra and signal processing. Taking a computational approach, this book explains how ANNs provide solutions in real time, and allow the visualization and development of new techniques and architectures. Features include: * A guide to the fundamental mathematics of neurocomputing. * A review of neural network models and an analysis of their associated algorithms. * State-of-the-art procedures to solve optimization problems. * Computer simulation programs MATLAB, TUTSIM and SPICE illustrate the validity and performance of the algorithms and architectures described. The authors encourage the reader to be creative in visualizing new approaches and detail how

other specialized computer programs can evaluate performance. * Each chapter concludes with a short bibliography. * Illustrative worked examples, questions and problems assist self study. The authors' self-contained approach will appeal to a wide range of readers, including professional engineers working in computing, optimization, operational research, systems identification and control theory. Undergraduate and postgraduate students in computer science, electrical and electronic engineering will also find this text invaluable. In particular, the text will be ideal to supplement courses in circuit analysis and design, adaptive systems, control systems, signal processing and parallel computing. B.G. Teubner Stuttgart

AI is already a massive presence in our lives. Be it in the learning algorithms that suggest products on Amazon, intelligent thermostats and voice recognition from Nest, Siri and Alexa. How does the fact that learning -and quasi intelligent- machines are becoming part of the process of construction, maintenance and everyday life change our habits towards the built environment? How can AI contribute to the experience of architecture in this environment? Is it part of the materialization of architecture?

This Workshop focuses on such issues as control algorithms which are suitable for real-time use, computer architectures which are suitable for real-time control algorithms, and applications for real-time control issues in the areas of parallel algorithms, multiprocessor systems, neural networks, fault-tolerance systems, real-time robot control identification, real-time filtering

algorithms, control algorithms, fuzzy control, adaptive and self-tuning control, and real-time control applications.

" The topic of this book the creation of software programs displaying broad, deep, human-style general intelligence is a grand and ambitious one. And yet it is far from a frivolous one: what the papers in this publication illustrate is that it is a fit and proper subject for serious science and engineering exploration. No one has yet created a software program with human-style or (even roughly) human-level general intelligence but we now have a sufficiently rich intellectual toolkit that it is possible to think about such a possibility in detail, and make serious attempts at design, analysis and engineering. possibility in detail, and make serious attempts at design, analysis and engineering. This is the situation that led to the organization of the 2006 AGIRI (Artificial General Intelligence Research Institute) workshop; and to the decision to publish a book from contributions by the speakers at the conference. The material presented here only scratches the surface of the AGI-related R&D work that is occurring around the world at this moment. But the editors are pleased to have had the chance to be involved in organizing and presenting at least a small percentage of the contemporary progress. "

This book constitutes the refereed proceedings of the 8th IAPR TC3 International Workshop on Artificial Neural Networks in Pattern Recognition, ANNPR 2018, held in Siena, Italy, in September 2018. The 29 revised full papers presented together with 2 invited papers were carefully reviewed and selected from 35 submissions. The papers present and discuss the latest research in all areas of neural network- and machine learning-based pattern recognition. They are organized in two sections: learning algorithms and architectures, and applications. Chapter "Bounded Rational

Decision-Making with Adaptive Neural Network Priors" is available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

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 four volume set LNCS 9528, 9529, 9530 and 9531 constitutes the refereed proceedings of the 15th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2015, held in Zhangjiajie, China, in November 2015. The 219 revised full papers presented together with 77 workshop papers in these four volumes were carefully reviewed and selected from 807 submissions (602 full papers and 205 workshop papers). The first volume comprises the following topics: parallel and distributed architectures; distributed and network-based computing and internet of things and cyber-physical-social computing. The second volume comprises topics such as big data and its applications and parallel and distributed algorithms. The topics of the third volume are: applications of parallel and distributed computing and service dependability and security in distributed and parallel systems. The covered topics of the fourth volume are: software systems and programming models and performance modeling and evaluation.
Providing the most comprehensive source available, this book surveys the state of the art in artificial intelligence (AI) as it relates to architecture. This book is organized in four parts: theoretical foundations, tools and techniques, AI in research, and AI in architectural practice. It provides a framework for the issues surrounding AI and offers a variety of perspectives. It contains 24 consistently illustrated contributions examining seminal work on AI from around the world, including the United States, Europe, and Asia. It articulates current theoretical and practical methods, offers critical views on tools and techniques, and suggests future directions for meaningful

uses of AI technology. Architects and educators who are concerned with the advent of AI and its ramifications for the design industry will find this book an essential reference. Increase the performance of various neural network architectures using NEAT, HyperNEAT, ES-HyperNEAT, Novelty Search, SAFE, and deep neuroevolution Key Features Implement neuroevolution algorithms to improve the performance of neural network architectures Understand evolutionary algorithms and neuroevolution methods with real-world examples Learn essential neuroevolution concepts and how they are used in domains including games, robotics, and simulations Book Description Neuroevolution is a form of artificial intelligence learning that uses evolutionary algorithms to simplify the process of solving complex tasks in domains such as games, robotics, and the simulation of natural processes. This book will give you comprehensive insights into essential neuroevolution concepts and equip you with the skills you need to apply neuroevolution-based algorithms to solve practical, real-world problems. You'll start with learning the key neuroevolution concepts and methods by writing code with Python. You'll also get hands-on experience with popular Python libraries and cover examples of classical reinforcement learning, path planning for autonomous agents, and developing agents to autonomously play Atari games. Next, you'll learn to solve common and not-so-common challenges in natural computing using neuroevolution-based algorithms. Later, you'll understand how to apply neuroevolution strategies to existing neural network designs to improve training and inference performance. Finally, you'll gain clear insights into the topology of neural networks and how neuroevolution allows you to develop complex networks, starting with simple ones. By the end of this book, you will not only have explored existing neuroevolution-based algorithms, but also have the skills you need to apply them in your

research and work assignments. What you will learn Discover the most popular neuroevolution algorithms – NEAT, HyperNEAT, and ES-HyperNEAT Explore how to implement neuroevolution-based algorithms in Python Get up to speed with advanced visualization tools to examine evolved neural network graphs Understand how to examine the results of experiments and analyze algorithm performance Delve into neuroevolution techniques to improve the performance of existing methods Apply deep neuroevolution to develop agents for playing Atari games Who this book is for This book is for machine learning practitioners, deep learning researchers, and AI enthusiasts who are looking to implement neuroevolution algorithms from scratch. Working knowledge of the Python programming language and basic knowledge of deep learning and neural networks are mandatory.

This three-volume set LNCS 12452, 12453, and 12454 constitutes the proceedings of the 20th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2020, in New York City, NY, USA, in October 2020. The total of 142 full papers and 5 short papers included in this proceedings volumes was carefully reviewed and selected from 495 submissions. ICA3PP is covering the many dimensions of parallel algorithms and architectures, encompassing fundamental theoretical approaches, practical experimental projects, and commercial components and systems. As applications of computing systems have permeated in every aspects of daily life, the power of computing system has become increasingly critical. This conference provides a forum for academics and practitioners from countries around the world to exchange ideas for improving the efficiency, performance, reliability, security and interoperability of computing systems and applications. ICA3PP 2020 focus on two broad areas of parallel and distributed computing, i.e. architectures, algorithms and

networks, and systems and applications.

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.

The four-volume set LNCS 11334-11337 constitutes the proceedings of the 18th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2018, held in Guangzhou, China, in November 2018. The 141 full and 50 short papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on Distributed and Parallel Computing; High Performance Computing; Big Data and Information Processing; Internet of Things and Cloud Computing; and Security and Privacy in Computing.

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.
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

COMMUNICATION NETWORKS AND SERVICE MANAGEMENT IN THE ERA OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING Discover the impact that new technologies are having on communication systems with this up-to-date and one-stop resource Communication Networks and Service Management in the Era of Artificial Intelligence and Machine Learning delivers a comprehensive overview of the impact of artificial intelligence (AI) and machine learning (ML) on service and network management. Beginning with a fulsome description of ML and AI, the book moves on to discuss management models, architectures, and frameworks. The authors also explore how AI and ML can be used in service management functions like the generation of workload profiles, service provisioning, and more. The book includes a handpicked selection of applications and case studies, as well as a

treatment of emerging technologies the authors predict could have a significant impact on network and service management in the future. Statistical analysis and data mining are also discussed, particularly with respect to how they allow for an improvement of the management and security of IT systems and networks. Readers will also enjoy topics like: A thorough introduction to network and service management, machine learning, and artificial intelligence An exploration of artificial intelligence and machine learning for management models, including autonomic management, policy-based management, intent based management, and network virtualization-based management Discussions of AI and ML for architectures and frameworks, including cloud systems, software defined networks, 5G and 6G networks, and Edge/Fog networks An examination of AI and ML for service management, including the automatic generation of workload profiles using unsupervised learning Perfect for information and communications technology educators, Communication Networks and Service Management in the Era of Artificial Intelligence and Machine Learning will also earn a place in the libraries of engineers and professionals who seek a structured reference on how the emergence of artificial intelligence and machine learning techniques is affecting service and network management.

Over the past few years, the demand for high speed Digital Signal Proces sing (DSP) has increased dramatically. New applications in real-time image processing, satellite communications, radar signal processing, pattern recogni tion, and real-time signal

detection and estimation require major improvements at several levels; algorithmic, architectural, and implementation. These perfor mance requirements can be achieved by employing parallel processing at all levels. Very Large Scale Integration (VLSI) technology supports and provides a good avenue for parallelism. Parallelism offers efficient sohitions to several problems which can arise in VLSI DSP architectures such as: 1. Intermediate data communication and routing: several DSP algorithms, such as FFT, involve excessive data routing and reordering. Parallelism is an efficient mechanism to minimize the silicon cost and speed up the pro cessing time of the intermediate middle stages. 2. Complex DSP applications: the required computation is almost doubled. Parallelism will allow two similar channels processing at the same time. The communication between the two channels has to be minimized. 3. Applicatilm specific systems: this emerging approach should achieve real-time performance in a cost-effective way. 4. Testability and fault tolerance: reliability has become a required feature in most of DSP systems. To achieve such property, the involved time overhead is significant. Parallelism may be the solution to maintain ac ceptable speed performance.