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

RONNIE HEATH

Automated Machine Learning Academic Press
" 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. "
Fundamentals of Neural Networks: Architectures, Algorithms and Applications IGI Global
Algorithms and Architectures of Artificial IntelligenceIOS Press
Algorithms and Architectures Springer Science & Business Media
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.
Introduction to Deep Learning IOS Press
Theoretical results suggest that in order to learn the kind of complicated 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.
Algorithms, Applications, Architectures of Artificial Neural Nets Springer
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
Hands-On Neuroevolution with Python CRC Press
"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."
Springer
This open access book presents the first comprehensive overview of general methods in Automated Machine Learning (AutoML), collects descriptions of existing systems based on these methods, and discusses the first series of international challenges of AutoML systems. The recent success of commercial ML applications and the rapid growth of the field has created a high demand for off-the-shelf ML methods that can be used easily and without expert knowledge. However, many of the recent machine learning successes crucially rely on human experts, who manually select appropriate ML architectures (deep learning architectures

or more traditional ML workflows) and their hyperparameters. To overcome this problem, the field of AutoML targets a progressive automation of machine learning, based on principles from optimization and machine learning itself. This book serves as a point of entry into this quickly-developing field for researchers and advanced students alike, as well as providing a reference for practitioners aiming to use AutoML in their work.
Algorithms and Architectures for Parallel Processing Algorithms and Architectures of Artificial Intelligence
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.
Analytical Formulation for Static and Temporal Patern Association Springer Science & Business Media
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 illustrated and highlight the introduced cross-layer design concepts.
Advances in Artificial General Intelligence Springer Nature
This book constitutes the refereed proceedings of the 6th IAPR TC3 International Workshop on Artificial Neural Networks in Pattern Recognition, ANNPR 2014, held in Montreal, QC, Canada, in October 2014. The 24 revised full papers presented were carefully reviewed and selected from 37 submissions for inclusion in this volume. They cover a large range of topics in the field of learning algorithms and architectures and discussing the latest research, results, and ideas in these areas.
Artificial Neural Networks in Pattern Recognition Pearson Education India
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.
Artificial Neural Nets and Genetic Algorithms Packt Publishing Ltd
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.
Machine Architectures for Artificial Intelligence Computing Elsevier
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.
An Introduction to AI for Architects Springer Science & Business Media
The three volume set LNCS 13155, 13156, and 13157 constitutes the refereed proceedings of the 21st International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2021, which was held online during December 3-5, 2021. The total of 145 full papers included in these proceedings were carefully reviewed and selected from 403 submissions. They cover the many dimensions of parallel algorithms and architectures including fundamental theoretical approaches, practical experimental projects, and commercial components and systems. The papers were organized in topical sections as follows: Part I, LNCS 13155: Deep learning models and applications; software systems and efficient algorithms; edge computing and edge intelligence; service dependability and security algorithms; data science; Part II, LNCS 13156: Software systems and efficient algorithms; parallel and distributed algorithms and applications; data science; edge computing and edge intelligence; blockchain systems; deept learning models and applications; IoT; Part III, LNCS 13157: Blockchain systems; data science; distributed and network-based computing; edge computing and edge intelligence; service dependability and security algorithms; software systems and efficient algorithms.
Learning Deep Architectures for AI Elsevier
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.
21st International Conference, ICA3PP 2021, Virtual Event, December 3–5, 2021, Proceedings, Part II Springer Science & Business Media
Artificial neural networks possess several properties that make them particularly attractive for applications to modelling and control of complex non-linear systems. Among these properties are their universal approximation ability, their parallel network structure and the availability of on- and off-line learning methods for the interconnection weights. However, dynamic models that contain neural network architectures might be highly non-linear and difficult to analyse as a result. Artificial Neural Networks for Modelling and Control of Non-Linear Systems investigates the subject from a system theoretical point of view. However the mathematical theory that is required from the reader is limited to matrix calculus, basic analysis, differential equations and basic linear system theory. No preliminary knowledge of neural networks is explicitly required. The book presents both classical and novel network architectures and learning algorithms for modelling and control. Topics include non-linear system identification, neural optimal control, top-down model based neural control design and stability analysis of neural control systems. A major contribution of this book is to introduce NLq Theory as an extension towards modern control theory, in order to analyze and synthesize non-linear systems that contain linear together with static non-linear operators that satisfy a sector condition: neural state space control systems are an example. Moreover, it turns out that NLq Theory is unifying with respect to many problems arising in neural networks, systems and control. Examples show that complex non-linear systems can be modelled and controlled within NLq theory, including mastering chaos. The didactic flavor of this book makes it suitable for use as a text for a course on Neural Networks. In addition, researchers and designers will find many important new techniques, in particular NLq emTheory, that have applications in control theory, system theory, circuit theory and Time Series Analysis.
Architectures, Algorithms, and Applications Springer
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
Design Challenges of Algorithm and Architecture Springer
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.
Architecture in the Age of Artificial Intelligence John Wiley & Sons
Artificial intelligence is everywhere – from the apps on our phones to the algorithms of search engines. Without us noticing, the AI revolution has arrived. But what does this mean for the world of design? The first volume in a two-book series, Architecture in the Age of Artificial Intelligence introduces AI for designers and considers its positive potential for the future of architecture and design. Explaining what AI is and how it works, the book examines how different manifestations of AI will impact the discipline and profession of architecture. Highlighting current case-studies as well as near-future applications, it shows how AI is already being

used as a powerful design tool, and how AI-driven information systems will soon transform the design of buildings and cities. Far-sighted, provocative and challenging, yet rooted in careful research and cautious speculation, this book, written by architect and theorist Neil Leach, is a must-read for all architects and designers – including students of architecture and all design professionals interested in keeping their practice at the cutting edge of technology.

Biologically Inspired Networking and Sensing: Algorithms and Architectures Createspace Independent Publishing Platform 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.

Related with Algorithms And Architectures Of Artificial Intelligence Frontiers In Artificial Intelligence And Applications:
• Cunningham Principles Of Environmental Science : click here