

Deep Belief Nets In C And Cuda C Volume 1 Restricted Boltzmann Machines And Supervised Feedforward Networks

Discover the essential building blocks of a common and powerful form of deep belief network: convolutional nets. This book shows you how the structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a 'thought process' that is capable of learning abstract concepts built from simpler primitives. These models are especially useful for image processing applications. At each step Deep Belief Nets in C++ and CUDA C: Volume 3 presents intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. Source code for all routines presented in the book, and the executable CONVNET program which implements these algorithms, are available for free download. What You Will Learn Discover convolutional nets and how to use them Build deep feedforward nets using locally connected layers, pooling layers, and softmax outputs Master the various programming algorithms required Carry out multi-threaded gradient computations and memory allocations for this threading Work with CUDA code implementations of all core computations, including layer activations and gradient calculations Make use of the CONVNET program and manual to explore convolutional nets and case studies Who This Book Is For Those who have at least a basic knowledge of neural networks and some prior programming experience, although some C++ and CUDA C is recommended.

The five volume set LNCS 7663, LNCS 7664, LNCS 7665, LNCS 7666 and LNCS 7667 constitutes the proceedings of the 19th International Conference on Neural Information Processing, ICONIP 2012, held in Doha, Qatar, in November 2012. The 423 regular session papers presented were carefully reviewed and selected from numerous submissions. These papers cover all major topics of theoretical research, empirical study and applications of neural information processing research. The 5 volumes represent 5 topical sections containing articles on theoretical analysis, neural modeling, algorithms, applications, as well as simulation and synthesis.

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.

Deep belief nets are one of the most exciting recent developments in artificial intelligence. The structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a 'thought process' that is capable of learning abstract concepts built from simpler primitives. A typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. This book presents the essential building blocks of the most common forms of deep belief nets. At each step the text provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. Source code for all routines presented in the book, and the DEEP program which implements these algorithms, are available for free download from the author's website.

The two volume set, LNCS 9886 + 9887, constitutes the proceedings of the 25th International Conference on Artificial Neural Networks, ICANN 2016, held in Barcelona, Spain, in September 2016. The 121 full papers included in this volume were carefully reviewed and selected from 227 submissions. They were organized in topical sections named: from neurons to networks; networks and dynamics; higher nervous functions; neuronal hardware; learning foundations; deep learning; classifications and forecasting; and recognition and navigation. There are 47 short paper abstracts that are included in the back matter of the volume.

This book is based on the International Conference on Decision Economics (DECON 2019). Highlighting the fact that important decision-making takes place in a range of critical subject areas and research fields, including economics, finance, information systems, psychology, small and international business, management, operations, and production, the book focuses on analytics as an emerging synthesis of sophisticated methodology and large data systems used to guide economic decision-making in an increasingly complex business environment. DECON 2019 was organised by the University of Chieti-Pescara (Italy), the National Chengchi University of Taipei (Taiwan), and the University of Salamanca (Spain), and was held at the Escuela politécnica Superior de Ávila, Spain, from 26th to 28th June, 2019. Sponsored by IEEE Systems Man and Cybernetics Society, Spain Section Chapter, and IEEE Spain Section (Technical Co-Sponsor), IBM, Indra, Viewnext, Global Exchange, AEPIA-and-APPIA, with the funding supporting of the Junta de Castilla y León, Spain (ID: SA267P18-Project co-financed with FEDER funds)

Build, test, and tune financial, insurance or other market trading systems using C++ algorithms and statistics. You've had an idea and have done some preliminary experiments, and it looks promising. Where do you go from here? Well, this book discusses and dissects this case study approach. Seemingly good backtest performance isn't enough to justify trading real money. You need to perform rigorous statistical tests of the system's validity. Then, if basic tests confirm the quality of your idea, you need to tune your system, not just for best performance, but also for robust behavior in the face of inevitable market changes. Next, you need to quantify its expected future behavior, assessing how bad its real-life performance might actually be, and whether you can live with that. Finally, you need to find its theoretical performance limits so you know if its actual trades conform to this theoretical expectation, enabling you to dump the system if it does not live up to expectations. This book does not contain any sure-fire, guaranteed-riches trading systems.

Those are a dime a dozen... But if you have a trading system, this book will provide you with a set of tools that will help you evaluate the potential value of your system, tweak it to improve its profitability, and monitor its on-going performance to detect deterioration before it fails catastrophically. Any serious market trader would do well to employ the methods described in this book. What You Will Learn See how the 'spaghetti-on-the-wall' approach to trading system development can be done legitimately Detect overfitting early in development Estimate the probability that your system's backtest results could have been due to just good luck Regularize a predictive model so it automatically selects an optimal subset of indicator candidates Rapidly find the global optimum for any type of parameterized trading system Assess the ruggedness of your trading system against market changes Enhance the stationarity and information content of your proprietary indicators Nest one layer of walkforward analysis inside another layer to account for selection bias in complex trading systems Compute a lower bound on your system's mean future performance Bound expected periodic returns to detect on-going system deterioration before it becomes severe Estimate the probability of catastrophic drawdown Who This Book Is For Experienced C++ programmers, developers, and software engineers. Prior experience with rigorous statistical procedures to evaluate and maximize the quality of systems is recommended as well.

This three volume set, CCIS 771, 772, 773, constitutes the refereed proceedings of the CCF Chinese Conference on Computer Vision, CCCV 2017, held in Tianjin, China, in October 2017. The total of 174 revised full papers presented in three volumes were carefully reviewed and selected from 465 submissions. The papers are organized in the following topical sections: biological vision inspired visual method; biomedical image analysis; computer vision applications; deep neural network; face and posture analysis; image and video retrieval; image color and texture; image composition; image quality assessment and analysis; image restoration; image segmentation and classification; image-based modeling; object detection and classification; object identification; photography and video; robot vision; shape representation and matching; statistical methods and learning; video analysis and event recognition; visual salient detection.

Discover the essential building blocks of a common and powerful form of deep belief net: the autoencoder. You'll take this topic beyond current usage by extending it to the complex domain for signal and image processing applications. Deep Belief Nets in C++ and CUDA C: Volume 2 also covers several algorithms for preprocessing time series and image data. These algorithms focus on the creation of complex-domain predictors that are suitable for input to a complex-domain autoencoder. Finally, you'll learn a method for embedding class information in the input layer of a restricted Boltzmann machine. This facilitates generative display of samples from individual classes rather than the entire data distribution. The ability to see the features that the model has learned for each class separately can be invaluable. At each step this book provides you with intuitive motivation, a summary of the most important equations relevant to the topic, and highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. What You'll Learn Code for deep learning, neural networks, and AI using C++ and CUDA C Carry out signal preprocessing using simple transformations, Fourier transforms, Morlet wavelets, and more Use the Fourier Transform for image preprocessing Implement autoencoding via activation in the complex domain Work with algorithms for CUDA gradient computation Use the DEEP operating manual Who This Book Is For Those who have at least a basic knowledge of neural networks and some prior programming experience, although some C++ and CUDA C is recommended.

This comprehensive encyclopedia, in A-Z format, provides easy access to relevant information for those seeking entry into any aspect within the broad field of Machine Learning. Most of the entries in this preeminent work include useful literature references.

This proceedings book presents the latest research in the fields of information theory, communication system, computer science and signal processing, as well as other related technologies. Collecting selected papers from the 3rd Conference on Signal and Information Processing, Networking and Computers (ICSINC), held in Chongqing, China on September 13-15, 2017, it is of interest to professionals from academia and industry alike.

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.

Big Data Analytics examines large amounts of data to uncover hidden patterns, correlations and other insights. With today's technology, it's possible to analyze your data and get answers from it almost immediately - an effort that's slower and less efficient with more traditional business intelligence solutions. Deep learning (also known as deep structured learning, hierarchical learning or deep machine learning) is a branch of machine learning based on a set of algorithms that attempt to model high level abstractions in data.

Various deep learning architectures such as deep neural networks, convolutional deep neural networks, deep belief networks and recurrent neural networks have been applied to fields like computer vision, automatic speech recognition, natural language processing, audio recognition and bioinformatics where they have been shown to produce state-of-the-art results on various tasks. Deep learning has been characterized as a buzzword, or a rebranding of neural networks. This book delves in big data and deep learning techniques

This book is a printed edition of the Special Issue "Short-Term Load Forecasting by Artificial Intelligent Technologies" that was published in Energies

Neuromorphic engineering has just reached its 25th year as a discipline. In the first two decades neuromorphic engineers focused on building models of sensors, such as silicon cochleas and retinas, and building blocks such as silicon neurons and synapses. These designs have honed our skills in implementing sensors and neural networks in VLSI using analog and mixed mode circuits. Over the last decade the address event representation has been used to interface devices and computers from different designers and even different groups. This facility has been essential for our ability to combine sensors, neural networks, and actuators into neuromorphic systems. More recently, several big projects have emerged to build very large scale neuromorphic systems. The Telluride Neuromorphic Engineering Workshop (since 1994) and the CapoCaccia Cognitive Neuromorphic Engineering Workshop (since 2009) have been instrumental not only in creating a strongly connected research community, but also in introducing different groups to each other's hardware. Many neuromorphic systems are first created at one of these workshops. With this special research topic, we showcase the state-of-the-art in neuromorphic systems.

Due to the growing use of web applications and communication devices, the use of data has increased throughout various industries. It is necessary to develop new techniques for managing data in order to ensure adequate usage. Deep learning, a subset of artificial intelligence and machine learning, has been recognized in various real-world applications such as computer vision, image processing, and pattern recognition. The deep learning approach has opened new opportunities that can make such real-life applications and tasks easier and more efficient. Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications is a vital reference source that trends in data analytics and potential technologies that will facilitate insight in various domains of science, industry, business, and consumer applications. It also explores the latest concepts, algorithms, and techniques of deep learning and data mining and analysis. Highlighting a range of topics such as natural language processing, predictive analytics, and deep neural networks, this multi-volume book is ideally designed for computer engineers, software developers, IT professionals, academicians, researchers, and upper-level students seeking current research on the latest trends in the field of deep learning.

This book presents a remarkable collection of chapters covering a wide range of topics in the areas of Computer Vision, both from theoretical and application perspectives. It gathers the proceedings of the Computer Vision Conference (CVC 2019), held in Las Vegas, USA from May 2 to 3, 2019. The conference attracted a total of 371 submissions from pioneering researchers, scientists, industrial engineers, and students all around the world. These submissions underwent a double-blind peer review process, after which 120 (including 7 poster papers) were selected for inclusion in these proceedings. The book's goal is to reflect the intellectual breadth and depth of current research on computer vision, from classical to intelligent scope. Accordingly, its respective chapters address state-of-the-art intelligent methods and techniques for solving real-world problems, while also outlining future research directions. Topic areas covered include Machine Vision and Learning, Data Science, Image Processing, Deep Learning, and Computer Vision Applications.

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

This book provides a comprehensive overview of deep learning (DL) in medical and healthcare applications, including the fundamentals and current advances in medical image analysis, state-of-the-art DL methods for medical image analysis and real-world, deep learning-based clinical computer-aided diagnosis systems. Deep learning (DL) is one of the key techniques of artificial intelligence (AI) and today plays an important role in numerous academic and industrial areas. DL involves using a neural network with many layers (deep structure) between input and output, and its main advantage of is that it can automatically learn data-driven, highly representative and hierarchical features and perform feature extraction and classification on one network. DL can be used to model or simulate an intelligent system or process using annotated training data. Recently, DL has become widely used in medical applications, such as anatomic modelling, tumour detection, disease classification, computer-aided diagnosis and surgical planning. This book is intended for computer science and engineering students and researchers, medical professionals and anyone interested using DL techniques.

In the current age of information explosion, newly invented technological sensors and software are now tightly integrated with our everyday lives. Many sensor processing algorithms have incorporated some forms of computational intelligence as part of their core framework in problem solving. These algorithms have the capacity to generalize and discover knowledge for themselves and learn new information whenever unseen data are captured. The primary aim of sensor processing is to develop techniques to interpret, understand, and act on information contained in the data. The interest of this book is in developing intelligent signal processing in order to pave the way for smart sensors. This involves mathematical advancement of nonlinear signal processing theory and its applications that extend far beyond traditional techniques. It bridges the boundary between

theory and application, developing novel theoretically inspired methodologies targeting both longstanding and emergent signal processing applications. The topic ranges from phishing detection to integration of terrestrial laser scanning, and from fault diagnosis to bio-inspired filtering. The book will appeal to established practitioners, along with researchers and students in the emerging field of smart sensors processing.

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.

Deep learning includes a subset of machine learning for processing the unsupervised data with artificial neural network functions. The major advantage of deep learning is to process big data analytics for better analysis and self-adaptive algorithms to handle more data. When applied to engineering, deep learning can have a great impact on the decision-making process. Deep Learning Applications and Intelligent Decision Making in Engineering is a pivotal reference source that provides practical applications of deep learning to improve decision-making methods and construct smart environments. Highlighting topics such as smart transportation, e-commerce, and cyber physical systems, this book is ideally designed for engineers, computer scientists, programmers, software engineers, research scholars, IT professionals, academicians, and postgraduate students seeking current research on the implementation of automation and deep learning in various engineering disciplines.

This book presents a compilation of selected papers from the first International Conference on Big Data Analysis and Deep Learning Applications (ICBDL 2018), and focuses on novel techniques in the fields of big data analysis, machine learning, system monitoring, image processing, conventional neural networks, communication, industrial information, and their applications. Readers will find insights to help them realize more efficient algorithms and systems used in real-life applications and contexts, making the book an essential reference guide for academic researchers, professionals, software engineers in the industry, and regulators of aviation authorities.

This authoritative resource presents a comprehensive illustration of modern Artificial Intelligence / Machine Learning (AI/ML) technology for radio frequency (RF) data exploitation. It identifies technical challenges, benefits, and directions of deep learning (DL) based object classification using radar data, including synthetic aperture radar (SAR) and high range resolution (HRR) radar. The performance of AI/ML algorithms is provided from an overview of machine learning (ML) theory that includes history, background primer, and examples. Radar data issues of collection, application, and examples for SAR/HRR data and communication signals analysis are discussed. In addition, this book presents practical considerations of deploying such techniques, including performance evaluation, energy-efficient computing, and the future unresolved issues.

Imagine looking at a picture on a printed page or computer screen, adjusting your eyes in a manner that most people can learn easily, and suddenly having objects pop out at you in vivid 3D. Many people have already experienced this with the Magic Eye and related posters that were massively popular in the 90's. But what is not so well known is that this single-image stereogram technology has come a long way since those early days. The big breakthrough came when algorithms were discovered that could map textures onto the surface of single-image stereograms. I believe this is the only available book that delves deeply into stereogram algorithms, including highly documented C++ source code. These algorithms can be used to great effect by artists to create works of art that are far beyond the crude stereograms of yesteryear. Perhaps even more importantly, the ability to display depth maps in clear stereo using only a single printed image can be invaluable for scientific presentations. This book is an essential resource for anyone writing programs for stereogram generation.

This book constitutes - in conjunction with the two-volume set LNCS 10954 and LNCS 10955 - the refereed proceedings of the 14th International Conference on Intelligent Computing, ICIC 2018, held in Wuhan, China, in August 2018. The 275 full papers and 72 short papers of the three proceedings volumes were carefully reviewed and selected from 632 submissions. The papers are organized in topical sections such as Evolutionary Computation and Learning; Neural Networks; Pattern Recognition; Image Processing; Information Security; Virtual Reality and Human-Computer Interaction; Business Intelligence and Multimedia Technology; Biomedical Informatics Theory and Methods; Swarm Intelligence and Optimization; Natural Computing; Quantum Computing; Intelligent Computing in Computer Vision; Fuzzy Theory and Algorithms; Machine Learning; Systems Biology; Intelligent Systems and Applications for Bioengineering; Evolutionary Optimization: Foundations and Its Applications to Intelligent Data Analytics; Swarm Evolutionary Algorithms for Scheduling and Combinatorial Optimization; Swarm Intelligence and Applications in Combinatorial Optimization; Advances in Metaheuristic Optimization Algorithm; Advances in Image Processing and Pattern Techniques; Bioinformatics.

Deep belief nets are one of the most exciting recent developments in artificial intelligence. The structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a 'thought process' that is capable of learning abstract concepts built from simpler primitives. A typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. This book presents the essential building blocks of a common and powerful form of deep belief net: convolutional nets. These models are especially useful for image processing applications. At each step the text provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. Source code for all routines presented in the book, and the executable CONVNET program which implements these algorithms, are available for free download from the author's website. Source code for the complete CONVNET program is not available, as much of it is highly specialized Windows interface code. Readers are responsible for writing their own main program, with all interface routines. You may freely use all of the core convolutional net routines in this book, as long as you remember that it is experimental code that comes with absolutely no guaranty of correct operation.

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.

Deep Belief Nets in C++ and CUDA C: Volume 2 Autoencoding in the Complex Domain Apress

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.

Discover the essential building blocks of the most common forms of deep belief networks. At each step this book provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. The first of three in a series on C++ and CUDA C deep learning and belief nets, *Deep Belief Nets in C++ and CUDA C: Volume 1* shows you how the structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a thought process that is capable of learning abstract concepts built from simpler primitives. As such, you'll see that a typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. All the routines and algorithms presented in the book are available in the code download, which also contains some libraries of related routines. What You Will Learn Employ deep learning using C++ and CUDA C Work with supervised feedforward networks Implement restricted Boltzmann machines Use generative samplings Discover why these are important Who This Book Is For Those who have at least a basic knowledge of neural networks and some prior programming experience, although some C++ and CUDA C is recommended.

This book is intended to present the state of the art in research on machine learning and big data analytics. The accepted chapters covered many themes including artificial intelligence and data mining applications, machine learning and applications, deep learning technology for big data analytics, and modeling, simulation, and security with big data. It is a valuable resource for researchers in the area of big data analytics and its applications.

This two-volume proceedings contains revised selected papers from the International Conference on Artificial Intelligence and Computational Intelligence, AICI 2010, held in Sanya, China, in October 2010.

The total of 105 high-quality papers presented were carefully reviewed and selected from 1216 submissions. The topics covered are: applications of artificial intelligence; automated problem solving; automatic programming; data mining and knowledge discovering; distributed AI and agents; expert and decision support systems; fuzzy logic and soft computing; intelligent information fusion; intelligent scheduling; intelligent signal processing; machine learning; machine vision; multi-agent systems; natural language processing; neural networks; pattern recognition; robotics; applications of computational intelligence; biomedical informatics and computation; fuzzy computation; genetic algorithms; immune computation; information security; intelligent agents and systems; nature computation; particle swarm optimization; and probabilistic reasoning.

Deep belief nets are one of the most exciting recent developments in artificial intelligence. The structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a 'thought process' that is capable of learning abstract concepts built from simpler primitives. A typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. This book presents the essential building blocks of a common and powerful form of deep belief net: the autoencoder. Volume II takes this topic beyond current usage by extending it to the complex domain, which is useful for many signal and image processing applications. Several algorithms for preprocessing time series and image data are also presented. These algorithms focus on the creation of complex-domain predictors that are suitable for input to a complex-domain autoencoder. Finally, this book provides a method for embedding class information in the input layer of a restricted Boltzmann machine. This facilitates generative display of samples from individual classes rather than the entire data distribution. The ability to see the features that the model has learned for each class separately can be invaluable. At each step the text provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. Source code for all routines presented in the book, and the DEEP program which implements these algorithms, are available for free download from the author's website.

Deep learning is a branch of machine learning based on a set of algorithms that attempt to model high level abstractions in data. In a simple case, there might be two sets of neurons: ones that receive an input signal and ones that send an output signal. When the input layer receives an input it passes on a modified version of the input to the next layer. In a deep network, there are many layers between the input and output (and the layers are not made of neurons but it can help to think of it that way), allowing the algorithm to use multiple processing layers, composed of multiple linear and non-linear transformations. Various deep learning architectures such as deep neural networks, convolutional deep neural networks, deep belief networks and recurrent neural networks have been applied to fields like computer vision, automatic speech recognition, natural language processing, audio recognition and bioinformatics where they have been shown to produce state-of-the-art results on various tasks.

BIG DATA ANALYTICS FOR INTERNET OF THINGS Discover the latest developments in IoT Big Data with a new resource from established and emerging leaders in the field Big Data Analytics for Internet of Things delivers a comprehensive overview of all aspects of big data analytics in Internet of Things (IoT) systems. The book includes discussions of the enabling technologies of IoT data analytics, types of IoT data analytics, challenges in IoT data analytics, demand for IoT data analytics, computing platforms, analytical tools, privacy, and security. The distinguished editors have included resources that address key techniques in the analysis of IoT data. The book demonstrates how to select the appropriate techniques to unearth valuable insights from IoT data and offers novel designs for IoT systems. With an abiding focus on practical strategies with concrete applications for data analysts and IoT professionals, Big Data Analytics for Internet of Things also offers readers: A thorough introduction to the Internet of Things, including IoT architectures, enabling technologies, and applications An exploration of the intersection between the Internet of Things and Big Data, including IoT as a source of Big Data, the unique characteristics of IoT data, etc. A discussion of the IoT data analytics, including the data analytical requirements of IoT data and the types of IoT analytics, including predictive, descriptive, and prescriptive analytics A treatment of machine learning techniques for IoT data analytics Perfect for professionals, industry practitioners, and researchers engaged in big data analytics related to IoT systems, Big Data Analytics for Internet of Things will also earn a place in the libraries of IoT designers and manufacturers interested in facilitating the efficient implementation of data analytics strategies.

This book constitutes the refereed proceedings of the 14th International Symposium on Neural Networks, ISNN 2017, held in Sapporo, Hakodate, and Muroran, Hokkaido, Japan, in June 2017. The 135 revised full papers presented in this two-volume set were carefully reviewed and selected from 259 submissions. The papers cover topics like perception, emotion and development, action and motor control, attractor and associative memory, neurodynamics, complex systems, and chaos.

In recent years, deep learning has fundamentally changed the landscapes of a number of areas in artificial intelligence, including speech, vision, natural language, robotics, and game playing. In particular, the striking success of deep learning in a wide variety of natural language processing (NLP) applications has served as a benchmark for the advances in one of the most important tasks in artificial intelligence.

This book reviews the state of the art of deep learning research and its successful applications to major NLP tasks, including speech recognition and understanding, dialogue systems, lexical analysis, parsing, knowledge graphs, machine translation, question answering, sentiment analysis, social computing, and natural language generation from images. Outlining and analyzing various research frontiers of NLP in the deep learning era, it features self-contained, comprehensive chapters written by leading researchers in the field. A glossary of technical terms and commonly used acronyms in the intersection of deep learning and NLP is also provided. The book appeals to advanced undergraduate and graduate students, post-doctoral researchers, lecturers and industrial researchers, as well as anyone interested in deep learning and natural language processing.

[Copyright: aa61c0c86e12f975913d84d30ac4fcca](#)