

Learning Deep Architectures For Ai

Recognizing the quirk ways to get this ebook **Learning Deep Architectures For Ai** is additionally useful. You have remained in right site to start getting this info. get the Learning Deep Architectures For Ai belong to that we offer here and check out the link.

You could purchase lead Learning Deep Architectures For Ai or acquire it as soon as feasible. You could speedily download this Learning Deep Architectures For Ai after getting deal. So, next you require the ebook swiftly, you can straight acquire it. Its consequently very easy and suitably fats, isnt it? You have to favor to in this appearance

Research Handbook on Intellectual Property and Artificial Intelligence - Ryan Abbott 2022-12-13

This incisive Handbook offers novel theoretical and doctrinal insights alongside practical guidance on some of the most challenging issues in the field of artificial intelligence and intellectual property.

Featuring all original contributions from a diverse group of international thought leaders, including top academics, judges, regulators and eminent practitioners, it offers timely perspectives and research on the relationship of AI to copyright, trademark, design, patent and trade secret law.

Deep Learning Illustrated - Jon Krohn 2019-08-05

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of *Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning* Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. **Deep Learning Illustrated** is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by

Agláé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Machine Learning and Information Processing -

Debabala Swain 2021-04-02

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

Network and Parallel Computing - Ching-Hsien Hsu 2014-08-23

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

Advances in Deep Learning - M. Arif Wani 2019-03-14

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

models.

Deep Learning and the Game of Go - Kevin Ferguson 2019-01-06

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

AND GAME AI Playing games with tree search
Getting started with neural networks Designing a
neural network for Go data Learning from data: a
deep-learning bot Deploying bots in the wild
Learning by practice: reinforcement learning
Reinforcement learning with policy gradients
Reinforcement learning with value methods
Reinforcement learning with actor-critic methods
PART 3 - GREATER THAN THE SUM OF ITS
PARTS AlphaGo: Bringing it all together AlphaGo
Zero: Integrating tree search with reinforcement
learning

*Applications of Artificial Intelligence and Machine
Learning* - Ankur Choudhary 2021-07-27

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

Artificial Intelligence in the Age of Neural
Networks and Brain Computing - Robert Kozma
2018-10-30

Artificial Intelligence in the Age of Neural
Networks and Brain Computing demonstrates that
existing disruptive implications and applications of
AI is a development of the unique attributes of
neural networks, mainly machine learning,
distributed architectures, massive parallel
processing, black-box inference, intrinsic
nonlinearity and smart autonomous search engines.
The book covers the major basic ideas of brain-like
computing behind AI, provides a framework to

deep learning, and launches novel and intriguing
paradigms as future alternatives. The success of AI-
based commercial products proposed by top industry
leaders, such as Google, IBM, Microsoft, Intel and
Amazon can be interpreted using this book.

Developed from the 30th anniversary of the
International Neural Network Society (INNS) and
the 2017 International Joint Conference on Neural
Networks (IJCNN) Authored by top experts, global
field pioneers and researchers working on cutting-
edge applications in signal processing, speech
recognition, games, adaptive control and decision-
making Edited by high-level academics and
researchers in intelligent systems and neural
networks

Artificial Intelligence and Soft Computing - Leszek
Rutkowski 2017-06-01

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

The Principles of Deep Learning Theory - Daniel
A. Roberts 2022-05-26

This volume develops an effective theory approach
to understanding deep neural networks of practical
relevance.

Biologically Inspired Cognitive Architectures 2021 -
Valentin V. Klimov 2022-03-24

The book focuses on original approaches intended to
support the development of biologically inspired
cognitive architectures. It bridges together different
disciplines, including artificial intelligence,
linguistics, neuro- and social sciences, psychology
and philosophy of mind, among others. The chapters

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

Deep Learning - Li Deng 2014

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

Learning Deep Architectures for AI - Yoshua Bengio 2009

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.

Handbook on Neural Information Processing - Monica Bianchini 2013-04-12

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

Artificial Intelligence Applications in Specialty Crops - Yiannis Ampatzidis 2022-03-02

Artificial Intelligence for Sustainable Development: Theory, Practice and Future Applications - Aboul Ella Hassanien 2020-08-31

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 In-memory Architectures for Machine Learning - Mingu Kang 2020-01-30

This book describes the recent innovation of deep in-memory architectures for realizing AI systems that operate at the edge of energy-latency-accuracy

trade-offs. From first principles to lab prototypes, this book provides a comprehensive view of this emerging topic for both the practicing engineer in industry and the researcher in academia. The book is a journey into the exciting world of AI systems in hardware.

Artificial Intelligence, Machine Learning, and Deep Learning - Oswald Campesato 2020-01-23

This book begins with an introduction to AI, followed by machine learning, deep learning, NLP, and reinforcement learning. Readers will learn about machine learning classifiers such as logistic regression, k-NN, decision trees, random forests, and SVMs. Next, the book covers deep learning architectures such as CNNs, RNNs, LSTMs, and auto encoders. Keras-based code samples are included to supplement the theoretical discussion. In addition, this book contains appendices for Keras, TensorFlow 2, and Pandas. Features: Covers an introduction to programming concepts related to AI, machine learning, and deep learning Includes material on Keras, TensorFlow2 and Pandas

Multidisciplinary Approaches to Neural Computing
- Anna Esposito 2017-08-28

This book presents a collection of contributions in the field of Artificial Neural Networks (ANNs). The themes addressed are multidisciplinary in nature, and closely connected in their ultimate aim to identify features from dynamic realistic signal exchanges and invariant machine representations that can be exploited to improve the quality of life of their end users. Mathematical tools like ANNs are currently exploited in many scientific domains because of their solid theoretical background and effectiveness in providing solutions to many demanding tasks such as appropriately processing (both for extracting features and recognizing) mono- and bi-dimensional dynamic signals, solving strong nonlinearities in the data and providing general solutions for deep and fully connected architectures. Given the multidisciplinary nature of their use and the interdisciplinary characterization of the problems they are applied to – which range from

medicine to psychology, industrial and social robotics, computer vision, and signal processing (among many others) – ANNs may provide a basis for redefining the concept of information processing. These reflections are supported by theoretical models and applications presented in the chapters of this book. This book is of primary importance for: (a) the academic research community, (b) the ICT market, (c) PhD students and early-stage researchers, (d) schools, hospitals, rehabilitation and assisted-living centers, and (e) representatives of multimedia industries and standardization bodies.

Deep Learning - Ian Goodfellow 2016-11-10

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online

recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Hands-On Deep Learning Architectures with Python - Yuxi (Hayden) Liu 2019-04-30

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
Dynamic Fuzzy Machine Learning - Fanzhang Li 2017-12-04
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.

Neural Networks: Tricks of the Trade - Grégoire Montavon 2012-11-14

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

TensorFlow for Deep Learning - Bharath Ramsundar 2018-03-01

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with

experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

Memristors for Neuromorphic Circuits and Artificial Intelligence Applications - Jordi Suñé 2020-04-09

Artificial Intelligence (AI) has found many applications in the past decade due to the ever increasing computing power. Artificial Neural Networks are inspired in the brain structure and consist in the interconnection of artificial neurons through artificial synapses. Training these systems requires huge amounts of data and, after the network is trained, it can recognize unforeseen data and provide useful information. The so-called Spiking Neural Networks behave similarly to how the brain functions and are very energy efficient. Up to this moment, both spiking and conventional neural networks have been implemented in software programs running on conventional computing units. However, this approach requires high computing power, a large physical space and is energy inefficient. Thus, there is an increasing interest in developing AI tools directly implemented in hardware. The first hardware demonstrations have been based on CMOS circuits for neurons and specific communication protocols for synapses. However, to further increase training speed and energy efficiency while decreasing system size, the combination of CMOS neurons with memristor synapses is being explored. The

memristor is a resistor with memory which behaves similarly to biological synapses. This book explores the state-of-the-art of neuromorphic circuits implementing neural networks with memristors for AI applications.

Development and Analysis of Deep Learning

Architectures - Witold Pedrycz 2019-11-01

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.

Advances in Artificial Intelligence - Denilson

Barbosa 2015-04-28

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

AI 2016: Advances in Artificial Intelligence -

Byeong Ho Kang 2016-11-25

This book constitutes the refereed proceedings of the 29th Australasian Joint Conference on Artificial Intelligence, AI 2016, held in Hobart, TAS, Australia, in December 2016. The 40 full papers and 18 short papers presented together with 8 invited short papers were carefully reviewed and selected from 121 submissions. The papers are organized in topical sections on agents and multiagent systems; AI applications and innovations; big data; constraint

satisfaction, search and optimisation; knowledge representation and reasoning; machine learning and data mining; social intelligence; and text mining and NLP. The proceedings also contains 2 contributions of the AI 2016 doctoral consortium and 6 contributions of the SMA 2016.

Python Deep Learning - Ivan Vasilev 2019-01-16

Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries
Key Features
Build a strong foundation in neural networks and deep learning with Python libraries
Explore advanced deep learning techniques and their applications across computer vision and NLP
Learn how a computer can navigate in complex environments with reinforcement learning
Book Description
With the surge in artificial intelligence in applications catering to both business and consumer needs, deep learning is more important than ever for meeting current and future market demands. With this book, you'll explore deep learning, and learn how to put machine learning to use in your projects. This second edition of Python Deep Learning will get you up to speed with deep learning, deep neural networks, and how to train them with high-performance algorithms and popular Python frameworks. You'll uncover different neural network architectures, such as convolutional networks, recurrent neural networks, long short-term memory (LSTM) networks, and capsule networks. You'll also learn how to solve problems in the fields of computer vision, natural language processing (NLP), and speech recognition. You'll study generative model approaches such as variational autoencoders and Generative Adversarial Networks (GANs) to generate images. As you delve into newly evolved areas of reinforcement learning, you'll gain an understanding of state-of-the-art algorithms that are the main components behind popular games Go, Atari, and Dota. By the end of the book, you will be well-versed with the theory of deep learning along with its real-world applications. What you will learn
Grasp the

mathematical theory behind neural networks and deep learning processes Investigate and resolve computer vision challenges using convolutional networks and capsule networks Solve generative tasks using variational autoencoders and Generative Adversarial Networks Implement complex NLP tasks using recurrent networks (LSTM and GRU) and attention models Explore reinforcement learning and understand how agents behave in a complex environment Get up to date with applications of deep learning in autonomous vehicles Who this book is for This book is for data science practitioners, machine learning engineers, and those interested in deep learning who have a basic foundation in machine learning and some Python programming experience. A background in mathematics and conceptual understanding of calculus and statistics will help you gain maximum benefit from this book.

Deep Learning for Natural Language Processing -

Jason Brownlee 2017-11-21

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

Data Science - Qinglei Zhou 2018-09-10

This two volume set (CCIS 901 and 902) constitutes the refereed proceedings of the 4th International Conference of Pioneering Computer Scientists, Engineers and Educators, ICPCSEE 2018 (originally ICYCSEE) held in Zhengzhou, China, in September 2018. The 125 revised full papers presented in these two volumes were carefully reviewed and selected

from 1057 submissions. The papers cover a wide range of topics related to basic theory and techniques for data science including mathematical issues in data science, computational theory for data science, big data management and applications, data quality and data preparation, evaluation and measurement in data science, data visualization, big data mining and knowledge management, infrastructure for data science, machine learning for data science, data security and privacy, applications of data science, case study of data science, multimedia data management and analysis, data-driven scientific research, data-driven bioinformatics, data-driven healthcare, data-driven management, data-driven eGovernment, data-driven smart city/planet, data marketing and economics, social media and recommendation systems, data-driven security, data-driven business model innovation, social and/or organizational impacts of data science.

Efficient Processing of Deep Neural Networks -

Vivienne Sze 2020-06-24

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

throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas.

Big Data Analytics - Mrutyunjaya Panda 2018-12-12

Social networking has increased drastically in recent years, resulting in an increased amount of data being created daily. Furthermore, diversity of issues and complexity of the social networks pose a challenge in social network mining. Traditional algorithm software cannot deal with such complex and vast amounts of data, necessitating the development of novel analytic approaches and tools.

This reference work deals with social network aspects of big data analytics. It covers theory, practices and challenges in social networking. The book spans numerous disciplines like neural networking, deep learning, artificial intelligence, visualization, e-learning in higher education, e-healthcare, security and intrusion detection.

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

- Aboul-Ella Hassanien 2020-10-12

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

Deep Learning with PyTorch - Luca Pietro Giovanni Antiga 2020-07-01

“We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in

great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch

Key Features

- Written by PyTorch’s creator and key contributors
- Develop deep learning models in a familiar Pythonic way
- Use PyTorch to build an image classifier for cancer detection
- Diagnose problems with your neural network and improve training with data augmentation

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About The Book

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It’s great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you’ll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks.

What You Will Learn

- Understanding deep learning data structures such as tensors and neural networks
- Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results
- Implementing modules and loss functions
- Utilizing pretrained models from PyTorch Hub
- Methods for training networks with limited inputs
- Sifting through unreliable results to diagnose and fix problems in your neural network
- Improve your results with augmented data, better model architecture, and fine tuning

This Book Is Written For

- Python programmers with an interest in machine learning.
- No experience with PyTorch or other deep learning frameworks is

required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1

Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Deep Learning: Concepts and Architectures - Witold Pedrycz 2019-10-29

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

detailed descriptions of the current trends in the design and analysis of deep learning topologies, the book offers practical guidelines and presents competitive solutions to various areas of language modeling, graph representation, and forecasting.

Machine Learning - Marco Gori 2023-04-01

Machine Learning: A Constraint-Based Approach, Second Edition provides readers with a refreshing look at the basic models and algorithms of machine learning, with an emphasis on current topics of interest that include neural networks and kernel machines. The book presents the information in a truly unified manner that is based on the notion of learning from environmental constraints. It draws a path towards deep integration with machine learning that relies on the idea of adopting multivalued logic formalisms, such as in fuzzy systems. Special attention is given to deep learning, which nicely fits the constrained-based approach followed in this book. The book presents a simpler unified notion of regularization, which is strictly connected with the parsimony principle, including many solved exercises that are classified according to the Donald Knuth ranking of difficulty, which essentially consists of a mix of warm-up exercises that lead to deeper research problems. A software simulator is also included. Presents, in a unified manner, fundamental machine learning concepts, such as neural networks and kernel machines Provides in-depth coverage of unsupervised and semi-supervised learning, with new content in hot growth areas such as deep learning Includes a software simulator for kernel machines and learning from constraints that also covers exercises to facilitate learning Contains hundreds of solved examples and exercises chosen particularly for their progression of difficulty from simple to complex Supported by a free, downloadable companion book designed to facilitate students' acquisition of experimental skills

Demystifying AI for the Enterprise - Prashant Natarajan 2021-12-31

Artificial intelligence (AI) in its various forms --

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

Embedded Deep Learning - Bert Moons 2018-10-23

This book covers algorithmic and hardware implementation techniques to enable embedded

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

Neural Networks for Natural Language Processing - S., Sumathi 2019-11-29

Information in today's advancing world is rapidly expanding and becoming widely available. This eruption of data has made handling it a daunting and time-consuming task. Natural language processing (NLP) is a method that applies linguistics and algorithms to large amounts of this data to make it more valuable. NLP improves the interaction between humans and computers, yet there remains a lack of research that focuses on the practical implementations of this trending approach. Neural Networks for Natural Language Processing is a collection of innovative research on the methods and applications of linguistic information processing and its computational properties. This publication will support readers with performing sentence classification and language generation using neural networks, apply deep learning models to solve machine translation and conversation problems, and

apply deep structured semantic models on information retrieval and natural language applications. While highlighting topics including deep learning, query entity recognition, and information retrieval, this book is ideally designed for research and development professionals, IT

specialists, industrialists, technology developers, data analysts, data scientists, academics, researchers, and students seeking current research on the fundamental concepts and techniques of natural language processing.