

Convolutional Neural Networks For Sentence Classification

Deep Learning for Natural Language Processing
 Developing Networks using Artificial Intelligence
 Theory and Engineering of Dependable Computer Systems and Networks
 Deep Learning for Natural Language Processing
 Deep Learning for Computer Vision
 Deep Learning Essentials
 Representation Learning for Natural Language Processing
 Deep Learning-Based Approaches for Sentiment Analysis
 5G for Future Wireless Networks
 Natural Language Processing with PyTorch
 Advanced Data Mining and Applications
 Convolutional Neural Networks in Visual Computing
 Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications
 Natural Language Processing with TensorFlow
 Engineering Applications of Neural Networks
 Artificial Neural Networks and Machine Learning – ICANN 2022
 Natural Language Processing and Chinese Computing
 Neural Representations of Natural Language
 Tree-Based Convolutional Neural Networks
 Natural Language Processing and Chinese Computing
 Neural Network Tutorials - Herong's Tutorial Examples
 Neural Networks with Python
 Artificial Neural Networks and Machine Learning – ICANN 2017
 Authorship Attribution
 Neural Networks for Natural Language Processing
 Feed-Forward Neural Networks
 Deep Learning
 Computational Intelligence in Pattern Recognition
 Deep Learning Approaches to Text Production
 Web Information Systems and Applications
 Proceedings of International Conference on Computational Intelligence and Data Engineering
 Neural Network Methods in Natural Language Processing
 Chinese Lexical Semantics
 Algorithms and Architectures for Parallel Processing
 Deep Learning with Python
 Database Systems for Advanced Applications
 Guide to Convolutional Neural Networks
 Neural Network Methods for Natural Language Processing
 Strengthening Deep Neural Networks
 Supervised Machine Learning for Text Analysis in R

Convolutional Neural Networks For Sentence Classification

Downloaded from db.mwpai.edu by guest

BRADSHAW HOUSTON

Deep Learning for Natural Language Processing Springer

This book contains papers on selected aspects of dependability analysis in computer systems and networks, which were chosen for discussion during the 16th DepCoS-RELCOMEX conference held in Wrocław, Poland, from June 28 to July 2, 2021. Their collection will be a valuable source material for scientists, researchers, practitioners and students who are dealing with design, analysis and engineering of computer systems and networks and must ensure their dependable operation. Being probably the most complex technical systems ever engineered by man (and also—the most dynamically evolving ones), organization of contemporary computer systems cannot be interpreted only as structures built on the basis of (unreliable) technical resources. Their evaluation must take into account a specific blend of interacting people (their needs and behaviours), networks (together with mobile properties, cloud organization, Internet of Everything, etc.) and a large number of users dispersed geographically and constantly producing an unconceivable number of applications. Ever-growing number of research methods being continuously developed for dependability analyses apply the newest techniques of artificial and computational intelligence. Selection of papers in these proceedings illustrates diversity of multi-disciplinary topics which are considered in present-day dependability explorations.

Developing Networks using Artificial Intelligence IGI Global

This book constitutes the proceedings of the 16th International Conference on Advanced Data Mining and Applications, ADMA 2020, held in Foshan, China in November 2020. The 35 full papers presented together with 14 short papers were carefully reviewed and selected from 96 submissions. The papers were organized in topical sections named: Machine Learning; Text Mining; Graph Mining; Predictive Analytics; Recommender Systems; Privacy and Security; Query Processing; Data Mining Applications.

Theory and Engineering of Dependable Computer Systems and Networks Packt Publishing Ltd

This book constitutes the refereed proceedings of the 19th International Conference on Engineering Applications of Neural Networks, EANN 2019, held in Xersonisos, Crete, Greece, in May 2019. The 35 revised full papers and 5 revised short papers presented were carefully reviewed and selected from 72 submissions. The papers are organized in topical sections on AI in energy management - industrial applications; biomedical - bioinformatics modeling; classification - learning; deep learning; deep learning - convolutional ANN; fuzzy - vulnerability - navigation modeling; machine learning modeling - optimization; ML - DL financial modeling; security - anomaly detection; 1st PEINT workshop.

Deep Learning for Natural Language Processing O'Reilly Media

This must-read text/reference introduces the fundamental concepts of convolutional neural networks (ConvNets), offering practical guidance on using libraries to implement ConvNets in applications of traffic sign detection and classification. The work presents techniques for optimizing the

computational efficiency of ConvNets, as well as visualization techniques to better understand the underlying processes. The proposed models are also thoroughly evaluated from different perspectives, using exploratory and quantitative analysis. Topics and features: explains the fundamental concepts behind training linear classifiers and feature learning; discusses the wide range of loss functions for training binary and multi-class classifiers; illustrates how to derive ConvNets from fully connected neural networks, and reviews different techniques for evaluating neural networks; presents a practical library for implementing ConvNets, explaining how to use a Python interface for the library to create and assess neural networks; describes two real-world examples of the detection and classification of traffic signs using deep learning methods; examines a range of varied techniques for visualizing neural networks, using a Python interface; provides self-study exercises at the end of each chapter, in addition to a helpful glossary, with relevant Python scripts supplied at an associated website. This self-contained guide will benefit those who seek to both understand the theory behind deep learning, and to gain hands-on experience in implementing ConvNets in practice. As no prior background knowledge in the field is required to follow the material, the book is ideal for all students of computer vision and machine learning, and will also be of great interest to practitioners working on autonomous cars and advanced driver assistance systems.

Deep Learning for Computer Vision Packt Publishing Ltd

This book constitutes the refereed proceedings of the 15th International Conference on Web Information Systems and Applications, WISA 2018, held in Taiyuan, China, in September 2018. The 29 full papers presented together with 16 short papers were carefully reviewed and selected from 103 submissions. The papers cover topics such as machine learning and data mining; cloud computing and big data; information retrieval; natural language processing; data privacy and security; knowledge graphs and social networks; query processing; and recommendations.

Deep Learning Essentials Springer

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.

Representation Learning for Natural Language Processing Springer

This book features high-quality research papers presented at the 3rd International Conference on Computational Intelligence in Pattern Recognition (CIPR 2021), held at the Institute of Engineering and Management, Kolkata, West Bengal, India, on 24 - 25 April 2021. It includes practical development experiences in various areas of data analysis and pattern recognition, focusing on soft computing technologies, clustering and classification algorithms, rough set and fuzzy set theory, evolutionary computations, neural science and neural network systems, image processing, combinatorial pattern matching, social network analysis, audio and video data analysis, data mining in dynamic environments, bioinformatics, hybrid computing, big data analytics and deep learning. It also provides innovative solutions to the challenges in these areas and discusses recent developments.

Deep Learning-Based Approaches for Sentiment Analysis Springer

The two volume set, LNCS 10613 and 10614, constitutes the proceedings of then 26th International Conference on Artificial Neural Networks, ICANN 2017, held in Alghero, Italy, in September 2017. The 128 full papers included in this volume were carefully reviewed and selected from 270 submissions. They were organized in topical sections named: From Perception to Action; From Neurons to Networks; Brain Imaging; Recurrent Neural Networks; Neuromorphic Hardware; Brain Topology and Dynamics; Neural Networks Meet Natural and Environmental Sciences; Convolutional Neural Networks; Games and Strategy; Representation and Classification; Clustering; Learning from Data Streams and Time Series; Image Processing and Medical Applications; Advances in Machine Learning. There are 63 short paper abstracts that are included in the back matter of the volume.

5G for Future Wireless Networks Springer Nature

This two volume set LNCS 10177 and 10178 constitutes the refereed proceedings of the 22nd International Conference on Database Systems for Advanced Applications, DASFAA 2017, held in Suzhou, China, in March 2017. The 73 full papers, 9 industry papers, 4 demo papers and 3 tutorials were carefully selected from a total of 300 submissions. The papers are organized around the following topics: semantic web and knowledge management; indexing and distributed systems; network embedding; trajectory and time series data processing; data mining; query processing and optimization; text mining; recommendation; security, privacy, sensor and cloud; social network analytics; map matching and spatial keywords; query processing and optimization; search and information retrieval; string and sequence processing; stream data processing; graph and network data processing; spatial databases; real time data processing; big data; social networks and graphs.

Natural Language Processing with PyTorch Springer

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. *Deep Learning for Natural Language Processing* follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of

IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Advanced Data Mining and Applications Machine Learning Mastery

This book constitutes the proceedings of the First International Conference on 5G for Future Wireless Networks, 5GWN 2017, held in Beijing, China, in April 2017. The 64 full papers were selected from 135 submissions and present the state of the art and practical applications of 5G technologies. The exponentially growing data traffic caused by the development of mobile Internet and smart phones requires powerful networks. The fifth generation (5G) techniques are promising to meet the requirements of this explosive data traffic in future mobile communications.

Convolutional Neural Networks in Visual Computing Springer

As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn't trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this adversarial input, and methods for increasing AI robustness to these attacks. If you're a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input Investigate methods used to generate adversarial input capable of fooling DNNs Explore real-world scenarios and model the adversarial threat Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data Examine some ways in which AI might become better at mimicking human perception in years to come

Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications Springer

Neural networks are a family of powerful machine learning models and this book focuses on their application to natural language data. The first half of the book (Parts I and II) covers the basics of supervised machine learning and feed-forward neural networks, the basics of working with machine learning over language data, and the use of vector-based rather than symbolic representations for words. It also covers the computation-graph abstraction, which allows to easily define and train arbitrary neural networks, and is the basis behind the design of contemporary neural network software libraries. The second part of the book (Parts III and IV) introduces more specialized neural network architectures, including 1D convolutional neural networks, recurrent neural networks, conditioned-generation models, and attention-based models. These architectures and techniques are the driving force behind state-of-the-art algorithms for machine translation, syntactic parsing, and many other applications. Finally, we also discuss tree-shaped networks, structured prediction, and the prospects of multi-task learning.

Natural Language Processing with TensorFlow Springer Nature

This book offers an introduction to modern natural language processing using machine learning, focusing on how neural networks create a machine interpretable representation of the meaning of natural language. Language is crucially linked to ideas – as Webster's 1923 "English Composition and Literature" puts it: "A sentence is a group of words expressing a complete thought". Thus the representation of sentences and the words that make them up is vital in advancing artificial intelligence and other "smart" systems currently being developed. Providing an overview of the research in the area, from Bengio et al.'s seminal work on a "Neural Probabilistic Language Model" in 2003, to the latest techniques, this book enables readers to gain an understanding of how the techniques are related and what is best for their purposes. As well as an introduction to neural networks in general and recurrent neural networks in particular, this book details the methods used for representing words, senses of words, and larger structures such as sentences or documents. The book highlights practical implementations and discusses many aspects that are often overlooked or misunderstood. The book includes thorough instruction on challenging areas such as hierarchical softmax and negative sampling, to ensure the reader fully and easily understands the details of how the algorithms function. Combining practical aspects with a more traditional review of the literature, it is directly applicable to a broad readership. It is an invaluable introduction for early graduate students working in natural language processing; a trustworthy guide for industry developers wishing to make use of recent innovations; and a sturdy bridge for researchers already familiar with linguistics or machine learning wishing to understand the other.

Engineering Applications of Neural Networks "O'Reilly Media, Inc."

This open access book provides an overview of the recent advances in representation learning theory, algorithms and applications for natural language processing (NLP). It is divided into three parts. Part I presents the representation learning techniques for multiple language entries, including words, phrases, sentences and documents. Part II then introduces the representation techniques for those objects that are closely related to NLP, including entity-based world knowledge, sememe-based linguistic knowledge, networks, and cross-modal entries. Lastly, Part III provides open resource tools for representation learning techniques, and discusses the remaining challenges and future research directions. The theories and algorithms of representation learning presented can also benefit other related domains such as machine learning, social network analysis, semantic Web, information retrieval, data mining and computational biology. This book is intended for advanced undergraduate and graduate students, post-doctoral fellows, researchers, lecturers, and industrial engineers, as well as anyone interested in representation learning and natural language processing.

Artificial Neural Networks and Machine Learning - ICANN 2022 Springer

This book mainly discusses the most important issues in artificial intelligence-aided future networks, such as applying different ML approaches to investigate solutions to intelligently monitor, control and optimize networking. The authors focus on four scenarios of successfully applying machine learning in network space. It also discusses the main challenge of network traffic intelligent awareness and introduces several machine learning-based traffic awareness algorithms, such as traffic classification, anomaly traffic identification and traffic prediction. The authors introduce some ML approaches like reinforcement learning to deal with network control problem in this book. Traditional works on the control plane largely rely on a manual process in configuring forwarding, which cannot be employed for today's network conditions. To address this issue, several artificial intelligence approaches for self-learning control strategies are introduced. In addition, resource management problems are ubiquitous in the

networking field, such as job scheduling, bitrate adaptation in video streaming and virtual machine placement in cloud computing. Compared with the traditional with-box approach, the authors present some ML methods to solve the complexity network resource allocation problems. Finally, semantic comprehension function is introduced to the network to understand the high-level business intent in this book. With Software-Defined Networking (SDN), Network Function Virtualization (NFV), 5th Generation Wireless Systems (5G) development, the global network is undergoing profound restructuring and transformation. However, with the improvement of the flexibility and scalability of the networks, as well as the ever-increasing complexity of networks, makes effective monitoring, overall control, and optimization of the network extremely difficult. Recently, adding intelligence to the control plane through AI&ML become a trend and a direction of network development This book's expected audience includes professors, researchers, scientists, practitioners, engineers, industry managers, and government research workers, who work in the fields of intelligent network. Advanced-level students studying computer science and electrical engineering will also find this book useful as a secondary textbook.

Natural Language Processing and Chinese Computing CRC Press

Write modern natural language processing applications using deep learning algorithms and TensorFlow Key Features Focuses on more efficient natural language processing using TensorFlow Covers NLP as a field in its own right to improve understanding for choosing TensorFlow tools and other deep learning approaches Provides choices for how to process and evaluate large unstructured text datasets Learn to apply the TensorFlow toolbox to specific tasks in the most interesting field in artificial intelligence Book Description Natural language processing (NLP) supplies the majority of data available to deep learning applications, while TensorFlow is the most important deep learning framework currently available. Natural Language Processing with TensorFlow brings TensorFlow and NLP together to give you invaluable tools to work with the immense volume of unstructured data in today's data streams, and apply these tools to specific NLP tasks. Thushan Ganegedara starts by giving you a grounding in NLP and TensorFlow basics. You'll then learn how to use Word2vec, including advanced extensions, to create word embeddings that turn sequences of words into vectors accessible to deep learning algorithms. Chapters on classical deep learning algorithms, like convolutional neural networks (CNN) and recurrent neural networks (RNN), demonstrate important NLP tasks as sentence classification and language generation. You will learn how to apply high-performance RNN models, like long short-term memory (LSTM) cells, to NLP tasks. You will also explore neural machine translation and implement a neural machine translator. After reading this book, you will gain an understanding of NLP and you'll have the skills to apply TensorFlow in deep learning NLP applications, and how to perform specific NLP tasks. What you will learn Core concepts of NLP and various approaches to natural language processing How to solve NLP tasks by applying TensorFlow functions to create neural networks Strategies to process large amounts of data into word representations that can be used by deep learning applications Techniques for performing sentence classification and language generation using CNNs and RNNs About employing state-of-the-art advanced RNNs, like long short-term memory, to solve complex text generation tasks How to write automatic translation programs and implement an actual neural machine translator from scratch The trends and innovations that are paving the future in NLP Who this book is for This book is for Python developers with a strong interest in deep learning, who want to learn how to leverage TensorFlow to simplify NLP tasks. Fundamental Python skills are assumed, as well as some knowledge of machine learning and undergraduate-level calculus and linear algebra. No previous natural language processing experience required, although some background in NLP or computational linguistics will be helpful.

Neural Representations of Natural Language Simon and Schuster

Best Sellers - Books :

- [The Woman In Me By Britney Spears](#)
- [The Mountain Is You: Transforming Self-sabotage Into Self-mastery By Brianna Wiest](#)
- [Jackie: Public, Private, Secret](#)
- [Twisted Lies \(twisted, 4\)](#)
- [Taylor Swift: A Little Golden Book Biography](#)
- [Killers Of The Flower Moon: The Osage Murders And The Birth Of The Fbi By David Grann](#)
- [The Five-star Weekend By Elin Hilderbrand](#)
- [Blowback: A Warning To Save Democracy From The Next Trump](#)
- [Harry Potter Paperback Box Set \(books 1-7\)](#)
- [How To Catch A Mermaid By Adam Wallace](#)

Step-by-step tutorials on deep learning neural networks for computer vision in python with Keras.

Tree-Based Convolutional Neural Networks Springer Nature

Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

Natural Language Processing and Chinese Computing Springer

Text data is important for many domains, from healthcare to marketing to the digital humanities, but specialized approaches are necessary to create features for machine learning from language. Supervised Machine Learning for Text Analysis in R explains how to preprocess text data for modeling, train models, and evaluate model performance using tools from the tidyverse and tidymodels ecosystem. Models like these can be used to make predictions for new observations, to understand what natural language features or characteristics contribute to differences in the output, and more. If you are already familiar with the basics of predictive modeling, use the comprehensive, detailed examples in this book to extend your skills to the domain of natural language processing. This book provides practical guidance and directly applicable knowledge for data scientists and analysts who want to integrate unstructured text data into their modeling pipelines. Learn how to use text data for both regression and classification tasks, and how to apply more straightforward algorithms like regularized regression or support vector machines as well as deep learning approaches. Natural language must be dramatically transformed to be ready for computation, so we explore typical text preprocessing and feature engineering steps like tokenization and word embeddings from the ground up. These steps influence model results in ways we can measure, both in terms of model metrics and other tangible consequences such as how fair or appropriate model results are.