

Supervised Sequence Labelling With Recurrent Neural Networks

This book constitutes the thoroughly refereed post-workshop proceedings of the 9th International Workshop on Statistical Atlases and Computational Models of the Heart: Atrial Segmentation and LV Quantification Challenges, STACOM 2018, held in conjunction with MICCAI 2018, in Granada, Spain, in September 2018. The 52 revised full workshop papers were carefully reviewed and selected from 60 submissions. The topics of the workshop included: cardiac imaging and image processing, machine learning applied to cardiac imaging and image analysis, atlas construction, statistical modelling of cardiac function across different patient populations, cardiac computational physiology, model customization, atlas based functional analysis, ontological schemata for data and results, integrated functional and structural analyses, as well as the pre-clinical and clinical applicability of these methods.

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

The proceedings set LNCS 11727, 11728, 11729, 11730, and 11731 constitute the proceedings of the 28th International Conference on Artificial Neural Networks, ICANN 2019, held in Munich, Germany, in September 2019. The total of 277 full papers and 43 short papers presented in these proceedings was carefully reviewed and selected from 494 submissions. They were organized in 5 volumes focusing on theoretical neural computation; deep learning; image processing; text and time series; and workshop and special sessions.

The three-volume set LNCS 10433, 10434, and 10435 constitutes the refereed proceedings of the 20th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI 2017, held in Quebec City, Canada, in September 2017. The 255 revised full papers presented were carefully reviewed and selected from 800 submissions in a two-phase review process. The papers have been organized in the following topical sections: Part I: atlas and surface-based techniques; shape and patch-based techniques; registration techniques, functional imaging, connectivity, and brain parcellation; diffusion magnetic resonance imaging (dMRI) and tensor/fiber processing; and image segmentation and modelling. Part II: optical imaging; airway and vessel analysis; motion and cardiac analysis; tumor processing; planning and simulation for medical interventions; interventional imaging and navigation; and medical image computing. Part III: feature extraction and classification techniques; and machine learning in medical image computing.

This book constitutes the refereed proceedings of the 12th IFIP WG 12.5 International

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

Conference on Artificial Intelligence Applications and Innovations, AIAI 2016, and three parallel workshops, held in Thessaloniki, Greece, in September 2016. The workshops are the Third Workshop on New Methods and Tools for Big Data, MT4BD 2016, the 5th Mining Humanistic Data Workshop, MHDW 2016, and the First Workshop on 5G - Putting Intelligence to the Network Edge, 5G-PINE 2016. The 30 revised full papers and 8 short papers presented at the main conference were carefully reviewed and selected from 65 submissions. The 17 revised full papers and 7 short papers presented at the 3 parallel workshops were selected from 33 submissions. The papers cover a broad range of topics such as artificial neural networks, classification, clustering, control systems - robotics, data mining, engineering application of AI, environmental applications of AI, feature reduction, filtering, financial-economics modeling, fuzzy logic, genetic algorithms, hybrid systems, image and video processing, medical AI applications, multi-agent systems, ontology, optimization, pattern recognition, support vector machines, text mining, and Web-social media data AI modeling.

This book constitutes the proceedings of the international workshops co-located with the 16th International Conference on Document Analysis and Recognition, ICDAR 2021, held in Lausanne, Switzerland, in September 2021. The total of 59 full and 12 short papers presented in this book were carefully selected from 96 contributions and divided into two volumes. Part I contains 29 full and 4 short papers that stem from the following meetings: ICDAR 2021 Workshop on Graphics Recognition (GREC); ICDAR 2021 Workshop on Camera-Based Document Analysis and Recognition (CBDAR); ICDAR 2021 Workshop on Arabic and Derived Script Analysis and Recognition (ASAR 2021); ICDAR 2021 Workshop on Computational Document Forensics (IWCDF). The main topics of the contributions are document processing; physical and logical layout analysis; text and symbol recognition; handwriting recognition; signature verification and document forensics, and others.

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source Deelearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular. Understand how deep networks evolved from neural network fundamentals. Explore the major deep network architectures, including Convolutional and Recurrent. Learn how to map specific deep networks to the right problem. Walk through the fundamentals of tuning general neural networks and specific deep network architectures. Use vectorization techniques for different data types with DataVec, DL4J's workflow tool. Learn how to use DL4J natively on Spark and Hadoop.

With a machine learning approach and less focus on linguistic details, this gentle introduction to natural language processing develops fundamental mathematical and deep learning models for NLP under a unified framework. NLP problems are systematically organised by their machine learning nature, including classification,

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

sequence labelling, and sequence-to-sequence problems. Topics covered include statistical machine learning and deep learning models, text classification and structured prediction models, generative and discriminative models, supervised and unsupervised learning with latent variables, neural networks, and transition-based methods. Rich connections are drawn between concepts throughout the book, equipping students with the tools needed to establish a deep understanding of NLP solutions, adapt existing models, and confidently develop innovative models of their own. Featuring a host of examples, intuition, and end of chapter exercises, plus sample code available as an online resource, this textbook is an invaluable tool for the upper undergraduate and graduate student.

This six volume set LNCS 11063 – 11068 constitutes the thoroughly refereed conference proceedings of the 4th International Conference on Cloud Computing and Security, ICCCS 2018, held in Haikou, China, in June 2018. The 386 full papers of these six volumes were carefully reviewed and selected from 1743 submissions. The papers cover ideas and achievements in the theory and practice of all areas of inventive systems which includes control, artificial intelligence, automation systems, computing systems, electrical and informative systems. The six volumes are arranged according to the subject areas as follows: cloud computing, cloud security, encryption, information hiding, IoT security, multimedia forensics

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. This book presents a systematic study of practices and theories for query understanding of search engines. These studies can be categorized into three major classes. The first class is to figure out what the searcher wants by extracting semantic meaning from the searcher's keywords, such as query classification, query tagging, and query intent understanding. The second class is to analyze search queries and then translate them into an enhanced query that can produce better search results, such as query spelling correction or query rewriting. The third class is to assist users in refining or suggesting queries in order to reduce users' search effort and satisfy their information needs, such as query auto-completion and query suggestion. Query understanding is a fundamental part of search engines. It is responsible to precisely infer the intent of the query formulated by the search user, to correct spelling errors in his/her query, to reformulate the query to capture its intent more accurately, and to guide the user in formulating a query with precise intent. The book will be invaluable to researchers and graduate students in computer or information science and specializing in information retrieval or web-based systems, as well as to researchers and programmers working on the development or improvement of products related to search engines.

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

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

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.

Natural language processing (NLP) went through a profound transformation in the mid-1980s when it shifted to make heavy use of corpora and data-driven techniques to analyze language. Since then, the use of statistical techniques in NLP has evolved in several ways. One such example of evolution took place in the late 1990s or early 2000s, when full-fledged Bayesian machinery was introduced to NLP. This Bayesian approach to NLP has come to accommodate various shortcomings in the frequentist approach and to enrich it, especially in the unsupervised setting, where statistical learning is done without target prediction examples. In this book, we cover the methods and algorithms that are needed to fluently read Bayesian learning papers in NLP and to do research in the area. These methods and algorithms are partially borrowed from both machine learning and statistics and are partially developed "in-house" in NLP. We cover inference techniques such as Markov chain Monte Carlo sampling and variational inference, Bayesian estimation, and nonparametric modeling. In response to rapid changes in the field, this second edition of the book includes a new chapter on representation learning and neural networks in the Bayesian context. We also cover fundamental concepts in Bayesian statistics such as prior distributions, conjugacy, and generative modeling. Finally, we review some of the fundamental modeling techniques in NLP, such as grammar modeling, neural networks and representation learning, and their use with Bayesian analysis.

This book presents a wealth of deep-learning algorithms and demonstrates their design process. It also highlights the need for a prudent alignment with the essential characteristics of the nature of learning encountered in the practical problems being tackled. Intended for readers interested in acquiring practical knowledge of analysis, design, and deployment of deep learning solutions to real-world problems, it covers a wide range of the paradigm's algorithms and their applications in diverse areas including imaging, seismic tomography, smart grids, surveillance and security, and health care, among others. Featuring systematic and comprehensive discussions on the development processes, their evaluation, and relevance, the book offers insights into fundamental design strategies for algorithms of deep learning.

This book constitutes the refereed proceedings of the 6th CCF International Conference on Natural Language Processing, NLPCC 2017, held in Dalian, China, in November 2017. The 47 full papers and 39 short papers presented were carefully reviewed and selected from 252 submissions. The papers are organized around the following topics: IR/search/bot; knowledge graph/IE/QA; machine learning; machine translation; NLP applications; NLP fundamentals; social networks; and text mining.

Spoken language understanding (SLU) is an emerging field in between speech and language processing, investigating human/ machine and human/ human communication by leveraging technologies from signal processing, pattern recognition, machine learning and artificial intelligence. SLU systems are designed to extract the meaning from speech utterances and its applications are vast, from voice search in mobile devices to meeting summarization, attracting interest from both commercial and academic sectors. Both human/machine and human/human communications can benefit from the application of SLU, using differing tasks and approaches to better understand and utilize such communications. This book covers the state-of-the-art approaches for the most popular SLU tasks with chapters written by well-known researchers in the respective fields. Key features include: Presents a fully integrated view of the two distinct disciplines of speech processing and language processing for SLU tasks. Defines what is possible today for SLU as an enabling technology for enterprise (e.g., customer care centers or company meetings), and consumer (e.g., entertainment, mobile, car, robot, or smart

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

environments) applications and outlines the key research areas. Provides a unique source of distilled information on methods for computer modeling of semantic information in human/machine and human/human conversations. This book can be successfully used for graduate courses in electronics engineering, computer science or computational linguistics. Moreover, technologists interested in processing spoken communications will find it a useful source of collated information of the topic drawn from the two distinct disciplines of speech processing and language processing under the new area of SLU.

It is our belief that researchers and practitioners acquire, through experience and word-of-mouth, techniques and heuristics that help them successfully apply neural networks to difficult real world problems. Often these "tricks" are theoretically well motivated. Sometimes they are the result of trial and error. However, their most common link is that they are usually hidden in people's heads or in the back pages of space-constrained conference papers. As a result newcomers to the field waste much time wondering why their networks train so slowly and perform so poorly. This book is an outgrowth of a 1996 NIPS workshop called Tricks of the Trade whose goal was to begin the process of gathering and documenting these tricks. The interest that the workshop generated motivated us to expand our collection and compile it into this book. Although we have no doubt that there are many tricks we have missed, we hope that what we have included will prove to be useful, particularly to those who are relatively new to the field. Each chapter contains one or more tricks presented by a given author (or authors). We have attempted to group related chapters into sections, though we recognize that the different sections are far from disjoint. Some of the chapters (e.g., 1, 13, 17) contain entire systems of tricks that are far more general than the category they have been placed in.

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

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

learning in their products or platforms. A website offers supplementary material for both readers and instructors.

This book constitutes the refereed proceedings of the 15th International Conference on Software Architecture, ECSA 2021, held in Sweden, in September 2021. Due to the COVID-19 pandemic, the conference was held virtually. For the Research Track, 11 full papers, presented together with 5 short papers, were carefully reviewed and selected from 58 submissions. The papers are organized in topical sections as follows: architectures for reconfigurable and self-adaptive systems; machine learning for software architecture; architectural knowledge, decisions, and rationale; architecting for quality attributes; architecture-centric source code analysis; and experiences and learnings from industrial case studies.

This book constitutes the refereed proceedings of the 14th China National Conference on Computational Linguistics, CCL 2014, and of the Third International Symposium on Natural Language Processing Based on Naturally Annotated Big Data, NLP-NABD 2015, held in Guangzhou, China, in November 2015. The 34 papers presented were carefully reviewed and selected from 283 submissions. The papers are organized in topical sections on lexical semantics and ontologies; semantics; sentiment analysis, opinion mining and text classification; machine translation; multilinguality in NLP; machine learning methods for NLP; knowledge graph and information extraction; discourse, coreference and pragmatics; information retrieval and question answering; social computing; NLP applications.

This book encompasses a collection of topics covering recent advances that are important to the Arabic language in areas of natural language processing, speech and image analysis. This book presents state-of-the-art reviews and fundamentals as well as applications and recent innovations. The book chapters by top researchers present basic concepts and challenges for the Arabic language in linguistic processing, handwritten recognition, document analysis, text classification and speech processing. In addition, it reports on selected applications in sentiment analysis, annotation, text summarization, speech and font analysis, word recognition and spotting and question answering. Moreover, it highlights and introduces some novel applications in vital areas for the Arabic language. The book is therefore a useful resource for young researchers who are interested in the Arabic language and are still developing their fundamentals and skills in this area. It is also interesting for scientists who wish to keep track of the most recent research directions and advances in this area.

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

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification,

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques. This book includes novel and state-of-the-art research discussions that articulate and report all research aspects, including theoretical and experimental prototypes and applications that incorporate sustainability into emerging applications. In recent years, sustainability and information and communication technologies (ICT) are highly intertwined, where sustainability resources and its management has attracted various researchers, stakeholders, and industrialists. The energy-efficient communication technologies have revolutionized the various smart applications like smart cities, healthcare, entertainment, and business. The book discusses and articulates emerging challenges in significantly reducing the energy consumption of communication systems and also explains development of a sustainable and energy-efficient mobile and wireless communication network. It includes best selected high-quality conference papers in different fields such as internet of things, cloud computing, data mining, artificial intelligence, machine learning, autonomous systems, deep learning, neural networks, renewable energy sources, sustainable wireless communication networks, QoS, network sustainability, and many other related areas.

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 to 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 the 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 topics range 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

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

sensor signal processing.

This open access book focuses on robot introspection, which has a direct impact on physical human-robot interaction and long-term autonomy, and which can benefit from autonomous anomaly monitoring and diagnosis, as well as anomaly recovery strategies. In robotics, the ability to reason, solve their own anomalies and proactively enrich owned knowledge is a direct way to improve autonomous behaviors. To this end, the authors start by considering the underlying pattern of multimodal observation during robot manipulation, which can effectively be modeled as a parametric hidden Markov model (HMM). They then adopt a nonparametric Bayesian approach in defining a prior using the hierarchical Dirichlet process (HDP) on the standard HMM parameters, known as the Hierarchical Dirichlet Process Hidden Markov Model (HDP-HMM). The HDP-HMM can examine an HMM with an unbounded number of possible states and allows flexibility in the complexity of the learned model and the development of reliable and scalable variational inference methods. This book is a valuable reference resource for researchers and designers in the field of robot learning and multimodal perception, as well as for senior undergraduate and graduate university students.

The two-volume set LNCS 9516 and 9517 constitutes the thoroughly refereed proceedings of the 22nd International Conference on Multimedia Modeling, MMM 2016, held in Miami, FL, USA, in January 2016. The 32 revised full papers and 52 poster papers were carefully reviewed and selected from 117 submissions. In addition 20 papers were accepted for five special sessions out of 38 submissions as well as 7 demonstrations (from 11 submissions) and 9 video showcase papers. The papers are organized in topical sections on video content analysis, social media analysis, object recognition and system, multimedia retrieval and ranking, multimedia representation, machine learning in multimedia, and interaction and mobile. The special sessions are: good practices in multimedia modeling; semantics discovery from multimedia big data; perception, aesthetics, and emotion in multimedia quality modeling; multimodal learning and computing for human activity understanding; and perspectives on multimedia analytics./div

In this textbook the author takes as inspiration recent breakthroughs in game playing to explain how and why deep reinforcement learning works. In particular he shows why two-person games of tactics and strategy fascinate scientists, programmers, and game enthusiasts and unite them in a common goal: to create artificial intelligence (AI). After an introduction to the core concepts, environment, and communities of intelligence and games, the book is organized into chapters on reinforcement learning, heuristic planning, adaptive sampling, function approximation, and self-play. The author takes a hands-on approach throughout, with Python code examples and exercises that help the reader understand how AI learns to play. He also supports the main text with detailed pointers to online machine learning frameworks, technical details for AlphaGo, notes on how to play and program Go and chess, and a comprehensive bibliography. The content is class-tested and suitable for advanced undergraduate and graduate courses on artificial intelligence and games. It's also appropriate for self-study by professionals engaged with applications of machine learning and with games development. Finally it's valuable for any reader engaged with the philosophical implications of artificial and general intelligence, games represent a modern Turing test of the power and limitations of AI.

Supervised Sequence Labelling with Recurrent Neural NetworksSpringer Science & Business Media

The eight-volume set comprising LNCS volumes 9905-9912 constitutes the refereed proceedings of the 14th European Conference on Computer Vision, ECCV 2016, held in Amsterdam, The Netherlands, in October 2016. The 415 revised papers presented were carefully reviewed and selected from 1480 submissions. The papers cover all aspects of computer vision and pattern recognition such as 3D computer vision; computational

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

photography, sensing and display; face and gesture; low-level vision and image processing; motion and tracking; optimization methods; physics-based vision, photometry and shape-from-X; recognition: detection, categorization, indexing, matching; segmentation, grouping and shape representation; statistical methods and learning; video: events, activities and surveillance; applications. They are organized in topical sections on detection, recognition and retrieval; scene understanding; optimization; image and video processing; learning; action activity and tracking; 3D; and 9 poster sessions.

Supervised sequence labelling is a vital area of machine learning, encompassing tasks such as speech, handwriting and gesture recognition, protein secondary structure prediction and part-of-speech tagging. Recurrent neural networks are powerful sequence learning tools—robust to input noise and distortion, able to exploit long-range contextual information—that would seem ideally suited to such problems. However their role in large-scale sequence labelling systems has so far been auxiliary. The goal of this book is a complete framework for classifying and transcribing sequential data with recurrent neural networks only. Three main innovations are introduced in order to realise this goal. Firstly, the connectionist temporal classification output layer allows the framework to be trained with unsegmented target sequences, such as phoneme-level speech transcriptions; this is in contrast to previous connectionist approaches, which were dependent on error-prone prior segmentation. Secondly, multidimensional recurrent neural networks extend the framework in a natural way to data with more than one spatio-temporal dimension, such as images and videos. Thirdly, the use of hierarchical subsampling makes it feasible to apply the framework to very large or high resolution sequences, such as raw audio or video. Experimental validation is provided by state-of-the-art results in speech and handwriting recognition.

The seven-volume set of LNCS 11301-11307, constitutes the proceedings of the 25th International Conference on Neural Information Processing, ICONIP 2018, held in Siem Reap, Cambodia, in December 2018. The 401 full papers presented were carefully reviewed and selected from 575 submissions. The papers address the emerging topics of theoretical research, empirical studies, and applications of neural information processing techniques across different domains. The first volume, LNCS 11301, is organized in topical sections on deep neural networks, convolutional neural networks, recurrent neural networks, and spiking neural networks.

This book constitutes the refereed proceedings of the 10th International Conference on Machine Learning and Data Mining in Pattern Recognition, MLDM 2014, held in St. Petersburg, Russia in July 2014. The 40 full papers presented were carefully reviewed and selected from 128 submissions. The topics range from theoretical topics for classification, clustering, association rule and pattern mining to specific data mining methods for the different multimedia data types such as image mining, text mining, video mining and Web mining.

The sixteen-volume set comprising the LNCS volumes 11205-11220 constitutes the refereed proceedings of the 15th European Conference on Computer Vision, ECCV 2018, held in Munich, Germany, in September 2018. The 776 revised papers presented were carefully reviewed and selected from 2439 submissions. The papers are organized in topical sections on learning for vision; computational photography; human analysis; human sensing; stereo and reconstruction; optimization; matching and recognition; video attention; and poster sessions. The eight-volume set comprising LNCS volumes 9905-9912 constitutes the refereed proceedings of the 14th European Conference on Computer Vision,

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

ECCV 2016, held in Amsterdam, The Netherlands, in October 2016. The 415 revised papers presented were carefully reviewed and selected from 1480 submissions. The papers cover all aspects of computer vision and pattern recognition such as 3D computer vision; computational photography, sensing and display; face and gesture; low-level vision and image processing; motion and tracking; optimization methods; physicsbased vision, photometry and shape-from-X; recognition: detection, categorization, indexing, matching; segmentation, grouping and shape representation; statistical methods and learning; video: events, activities and surveillance; applications. They are organized in topical sections on detection, recognition and retrieval; scene understanding; optimization; image and video processing; learning; action activity and tracking; 3D; and 9 poster sessions.

This three-volume set, LNAI 11670, LNAI 11671, and LNAI 11672 constitutes the thoroughly refereed proceedings of the 16th Pacific Rim Conference on Artificial Intelligence, PRICAI 2019, held in Cuvu, Yanuca Island, Fiji, in August 2019. The 111 full papers and 13 short papers presented in these volumes were carefully reviewed and selected from 265 submissions. PRICAI covers a wide range of topics such as AI theories, technologies and their applications in the areas of social and economic importance for countries in the Pacific Rim.

The purpose of this edited volume is to provide a comprehensive overview on the fundamentals of deep learning, introduce the widely-used learning architectures and algorithms, present its latest theoretical progress, discuss the most popular deep learning platforms and data sets, and describe how many deep learning methodologies have brought great breakthroughs in various applications of text, image, video, speech and audio processing. Deep learning (DL) has been widely considered as the next generation of machine learning methodology. DL attracts much attention and also achieves great success in pattern recognition, computer vision, data mining, and knowledge discovery due to its great capability in learning high-level abstract features from vast amount of data. This new book will not only attempt to provide a general roadmap or guidance to the current deep learning methodologies, but also present the challenges and envision new perspectives which may lead to further breakthroughs in this field. This book will serve as a useful reference for senior (undergraduate or graduate) students in computer science, statistics, electrical engineering, as well as others interested in studying or exploring the potential of exploiting deep learning algorithms. It will also be of special interest to researchers in the area of AI, pattern recognition, machine learning and related areas, alongside engineers interested in applying deep learning models in existing or new practical applications.

The Long Short-Term Memory network, or LSTM for short, is a type of recurrent neural network that achieves state-of-the-art results on challenging prediction problems. In this laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about LSTMs. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what

Online Library Supervised Sequence Labelling With Recurrent Neural Networks

LSTMs are, and how to develop a suite of LSTM models to get the most out of the method on your sequence prediction problems.

In today's modernized world, the field of healthcare has seen significant practical innovations with the implementation of computational intelligence approaches and soft computing methods. These two concepts present various solutions to complex scientific problems and imperfect data issues. This has made both very popular in the medical profession. There are still various areas to be studied and improved by these two schemes as healthcare practices continue to develop.

Computational Intelligence and Soft Computing Applications in Healthcare Management Science is an essential reference source that discusses the implementation of soft computing techniques and computational methods in the various components of healthcare, telemedicine, and public health. Featuring research on topics such as analytical modeling, neural networks, and fuzzy logic, this book is ideally designed for software engineers, information scientists, medical professionals, researchers, developers, educators, academicians, and students.

[Copyright: 6f6da7083fc94bf200ab1e12e54a25b9](#)