Convolutional Sequence To Sequence Learning Arxiv

These proceedings collect papers presented at the 11th International Conference on Multimedia & Network Information Systems (MISSI 2018), held from 12 to 14 September 2018 in Wroc?aw, Poland. The keynote lectures, given by four outstanding scientists, are also included here. The Conference attracted a great number of scientists from across Europe and beyond, and hosted the 6th International Workshop on Computational Intelligence for Multimedia Understanding as well as four special sessions. The majority of the papers describe various artificial intelligence (AI) methods applied to multimedia and natural language (NL) processing; they address hot topics such as virtual and augmented reality, identity recognition, video summarization, intelligent audio processing, accessing multilingual information and opinions, video games, and innovations in Web technologies. Accordingly, the proceedings provide a cutting-edge update on work being pursued in the rapidly evolving field of Multimedia and Internet Information Systems.

Recurrent Neural Networks (RNNs), which has the attractive properties of modelling sequences, has been dominant in speech field in the recent decades. Convolutional Neural Networks (CNNs) has been shown as an alternative to model sequences because of its capacity of reducing spectral variations and modeling spectral correlations in acoustic features for automatic speech recognition (ASR). Recent work suggests that complex numbers could be used as a richer feature representation than spectrum which may benefit the speech related tasks. In the thesis, we first cover the basic concepts in machine learning, building blocks of deep learning and discuss the popular methods that are capable of doing sequence-to-sequence modelling, specially convolutional neural networks, which is famous as a class of feed-forward nets. We then present two research work related to sequence-to-sequence modelling on speech. We introduce a new approach to address speech recognition with convolutional neural networks which shows the comparable results with their recurrent neural networks counterpart. In addition, we present a new model taking advantage of the representation in the complex domain and define complex convolutions, complex batch-normalization, complex weight initialization strategies. The new model results in state-of-the-art of speech spectrum prediction in a convolutional recurrent setting.

This book provides an overview of fake news detection, both through a variety of tutorial-style survey articles that capture advancements in the field from various facets and in a somewhat unique direction through expert perspectives from various disciplines. The approach is based on the idea that advancing the frontier on data science approaches for fake news is an interdisciplinary effort, and that perspectives from domain experts are crucial to shape the next generation of methods and tools. The fake news challenge cuts across a number of data science subfields such as graph analytics, mining of spatio-temporal data, information retrieval, natural language processing, computer vision and image processing, to name a few. This book will present a number of tutorial-style surveys that summarize a range of recent work in the field. In a unique feature, this book includes perspective notes from experts in disciplines such as linguistics, anthropology, medicine and politics that will help to shape the next generation of data science research in fake news. The main target groups of this book are academic and industrial researchers working in the area of data science, and with interests in devising and applying data science technologies for fake news detection. For young researchers such as PhD students, a review of data science work on fake news is provided, equipping them with enough know-how to start engaging in research within the area. For experienced researchers, the detailed descriptions of approaches will enable them to take seasoned choices in identifying promising directions for future research.

Sequential data is ubiquitous and modeling sequential data has been one of the most long-standing computer science problems. The goal of sequence modeling is to represent a sequence with a low-dimensional dense vector that incorporates as much information as possible. A fundamental type of information contained in sequences is the sequential dependency and a large body of research has been devoted to designing effective ways to capture it. Recently, sequence learning models such as recurrent neural networks (RNNs), temporal convolutional networks, and Transformer have gained tremendous popularity in modeling sequential data. Equipped with effective structures such as gating mechanisms, large receptive fields, and attention mechanisms, these models have achieved great success in many applications of a wide range of fields.However, besides the sequential dependency, sequences also exhibit side information that remains under-explored. Thus, in the thesis, we study the problem of sequence learning with side information. Specifically, we present our efforts devoted to building sequence learning models to effectively and efficiently capture side information that is commonly seen in sequential data. In addition, we show that side information can play an important role in sequence learning tasks as it can provide rich information that is complementary to the sequential dependency. More importantly, we apply our proposed models in various real-world applications and have achieved promising results.

This book constitutes the proceedings of the 18th Russian Conference on Artificial Intelligence, RCAI 2020, held in Moscow, Russia, in October 2020. The 27 full papers and 8 short papers presented in this volume were carefully reviewed and selected from 140 submissions. The conference deals with a wide range of topics, including data mining and knowledge discovery, text mining, reasoning, decisionmaking, natural language processing, vision, intelligent robotics, multi-agent systems,machine learning, AI in applied systems, and ontology engineering.

Deep Learning - 2 BOOK BUNDLE!! Deep Learning with Keras This book will introduce you to various supervised and unsupervised deep learning algorithms like the multilayer perceptron, linear regression and other more advanced deep convolutional and recurrent neural networks. You will also learn about image processing, handwritten recognition, object recognition and much more. Furthermore, you will get familiar with recurrent neural networks like LSTM and GAN as you explore processing sequence data like time series, text, and audio. The book will definitely be your best companion on this great deep learning journey with Keras introducing you to the basics you need to know in order to take next steps and learn more advanced deep neural networks. Here Is a Preview of What You'll Learn Here… The difference between deep learning and machine learning Deep neural networks Convolutional neural networks Building deep learning models with Keras Multi-layer perceptron network models Activation functions Handwritten recognition using MNIST Solving multi-class classification problems Recurrent neural networks and sequence classification And much more... Convolutional Neural Networks in Python This book covers the basics behind Convolutional Neural Networks by introducing you to this complex world of deep learning and artificial neural networks in a simple and easy to understand way. It is perfect for any beginner out there looking forward to learning more about this machine learning field. This book is all about how to use convolutional neural networks for various image, object and other common classification problems in Python. Here, we also take a deeper look into various Keras layer used for building CNNs we take a look at different activation functions and much more, which will eventually lead you to creating highly accurate models able of performing great task results on various image classification, object classification and other problems. Therefore, at the end of the book, you will have a better insight into this world, thus you will be more than prepared to deal with more complex and challenging tasks on your own. Here Is a Preview of What You'll Learn In This Book… Convolutional neural networks structure How convolutional neural networks actually work Convolutional neural networks applications The importance of convolution operator Different convolutional neural networks layers and their importance Arrangement of spatial parameters How and when to use stride and zero-padding Method of parameter sharing Matrix multiplication and its importance Pooling and dense layers Introducing non-linearity relu activation function How to train your convolutional neural network models using backpropagation How and why to apply dropout CNN model training process How to build a convolutional neural network Generating predictions and calculating loss functions How to train and evaluate your MNIST classifier How to build a simple image classification CNN And much, much more! Get this book bundle NOW and SAVE money!

This book collects 14 articles from the Special Issue entitled "Deep Learning Applications with Practical Measured Results in Electronics

Industries" of Electronics. Topics covered in this Issue include four main parts: (1) environmental information analyses and predictions, (2) unmanned aerial vehicle (UAV) and object tracking applications, (3) measurement and denoising techniques, and (4) recommendation systems and education systems. These authors used and improved deep learning techniques (e.g., ResNet (deep residual network), Faster-RCNN (faster regions with convolutional neural network), LSTM (long short term memory), ConvLSTM (convolutional LSTM), GAN (generative adversarial network), etc.) to analyze and denoise measured data in a variety of applications and services (e.g., wind speed prediction, air quality prediction, underground mine applications, neural audio caption, etc.). Several practical experiments were conducted, and the results indicate that the performance of the presented deep learning methods is improved compared with the performance of conventional machine learning methods.

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.

This book constitutes the proceedings of the 24th Annual Conference on Research in Computational Molecular Biology, RECOMB 2020, held in Padua, Italy, in May 2020. The 13 regular and 24 short papers presented were carefully reviewed and selected from 206 submissions. The papers report on original research in all areas of computational molecular biology and bioinformatics.

Sequence to Sequence Learning and Its Speech Applications

This proceedings book covers the theory, design and applications of computer networks, distributed computing and information systems. Today's networks are evolving rapidly, and there are several developing areas and applications. These include heterogeneous networking supported by recent technological advances in power wireless communications, along with silicon integration of various functionalities such as sensing, communications, intelligence and actuations, which is emerging as a critically important disruptive computer class based on a new platform, networking structure and interface that enables novel, low-cost and high-volume applications. However, implemeting these applications has sometimes been difficult due to interconnection problems. As such, different networks need to collaborate, and wired and next-generation wireless systems need to be integrated in order to develop high-performance computing solutions to address the problems arising from these networks' complexities. This ebook presents the latest research findings, as well as theoretical and practical perspectives on the innovative methods and development techniques related to the emerging areas of information networking and applications

The three volume proceedings LNAI 11051 – 11053 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2018, held in Dublin, Ireland, in September 2018. The total of 131 regular papers presented in part I and part II was carefully reviewed and selected from 535 submissions; there are 52 papers in the applied data science, nectar and demo track. The contributions were organized in topical sections named as follows: Part I: adversarial learning; anomaly and outlier detection; applications; classification; clustering and unsupervised learning; deep learningensemble methods; and evaluation. Part II: graphs; kernel methods; learning paradigms; matrix and tensor analysis; online and active learning; pattern and sequence mining; probabilistic models and statistical methods; recommender systems; and transfer learning. Part III: ADS data science applications; ADS e-commerce; ADS engineering and design; ADS financial and security; ADS health; ADS sensing and positioning; nectar track; and demo track.

This book constitutes the proceedings of the 10th Mexican Conference on Pattern Recognition, MCPR 2018, held in Puebla, Mexico, in June 2018. The 28 papers presented in this volume were carefully reviewed and selected from 44 submissions. They were organized in topical sections named: pattern recognition principles; deep learning, neural networks and associative memories; data mining; and computer vision.

In this thesis, the deep learning techniques called Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN) are used to address the problem of Optical Character Recognition (OCR). A special case of RNN called Long Short-Term Memory (LSTM) is used in this research to process the data sequentially. OCR is a process to convert the images containing characters into text. In this research, the images of the mathematical equations from Image-to-Latex 100K data set obtained from OPENAI organization is being used. The mathematical equations from the images are converted into Latex representation using deep learning techniques. The Latex texts were used to again recreate the mathematical equation to test the accuracy of the technique. Unlike previous techniques (Like INFTY) where models were fed with non-tokenized data, the proposed method used the tokenized data to be fed sequentially to the deep learning neural network. The sequential process helps the algorithms to keep track of the processed data and yield high accuracy. In this research, a new variant of LSTM called LSTM with peephole connections and Stochastic Hard Attention model was used. The performance of the proposed deep learning neural network, LSTM with peephole connections and Stochastic Hard Attention model is compared with INFTY (which uses no RNN) and WYGIWYS (which uses RNN). It has been found that the proposed algorithm gives a better accuracy of 76% as compared of 74% achieved by WYGIWYS.

The proceedings set LNCS 12891, LNCS 12892, LNCS 12893, LNCS 12894 and LNCS 12895 constitute the proceedings of the 30th International Conference on Artificial Neural Networks, ICANN 2021, held in Bratislava, Slovakia, in September 2021.* The total of 265 full papers presented in these proceedings was carefully reviewed and selected from 496 submissions, and organized in 5 volumes. In this volume, the papers focus on topics such as generative neural networks, graph neural networks, hierarchical and ensemble models, human pose estimation, image processing, image segmentation, knowledge distillation, and medical image processing. *The conference was held online 2021 due to the COVID-19 pandemic.

Learn how to build machine translation systems with deep learning from the ground up, from basic concepts to cutting-edge research.

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 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.

The LNCS 12115 constitutes the workshop papers which were held also online in conjunction with the 25th International Conference on Database Systems for Advanced Applications in September 2020. The complete conference includes 119 full papers presented together with 19 short papers plus 15 demo papers and 4 industrial papers in this volume were carefully reviewed and selected from a total of 487 submissions. DASFAA 2020 presents this year following five workshops: The 7th International Workshop on Big Data Management and Service (BDMS 2020) The 6th International Symposium on Semantic Computing and Personalization (SeCoP 2020) The 5th Big Data Quality Management (BDQM 2020) The 4th International Workshop on Graph Data Management and Analysis (GDMA 2020) The 1st International Workshop on Artificial Intelligence for Data Engineering (AIDE 2020)

This book provides readers with a comprehensive, state-of-the-art overview of approximate computing, enabling the design trade-off of accuracy for achieving better power/performance efficiencies, through the simplification of underlying computing resources. The authors describe in detail various efforts to generate approximate hardware systems, while still providing an overview of support techniques at other computing layers. The book is organized by techniques for various hardware components, from basic building blocks to general circuits and systems.

This book constitutes the refereed proceedings of the 7th Conference on Artificial Intelligence and Natural Language, AINL 2018, held in St. Petersburg, Russia, in October 2018. The 19 revised full papers were carefully reviewed and selected from 56 submissions and cover a wide range of topics, including morphology and word-level semantics, sentence and discourse representations, corpus linguistics, language resources, and social interaction analysis.

In recent years, deep learning has fundamentally changed the landscapes of a number of areas in artificial intelligence, including speech, vision, natural language, robotics, and game playing. In particular, the striking success of deep learning in a wide variety of natural language processing (NLP) applications has served as a benchmark for the advances in one of the most important tasks in artificial intelligence. This book reviews the state of the art of deep learning research and its successful applications to major NLP tasks, including speech recognition and understanding, dialogue systems, lexical analysis, parsing, knowledge graphs, machine translation, question answering, sentiment analysis, social computing, and natural language generation from images. Outlining and analyzing various research frontiers of NLP in the deep learning era, it features self-contained, comprehensive chapters written by leading researchers in the field. A glossary of technical terms and commonly used acronyms in the intersection of deep learning and NLP is also provided. The book appeals to advanced undergraduate and graduate students, post-doctoral researchers, lecturers and industrial researchers, as well as anyone interested in deep learning and natural language processing.

This two volume set of LNAI 11061 and LNAI 11062 constitutes the refereed proceedings of the 11th International Conference on Knowledge Science, Engineering and Management, KSEM 2018, held in Changchun, China, in August 2018. The 62 revised full papers and 26 short papers presented were carefully reviewed and selected from 262 submissions. The papers of the first volume are organized in the following topical sections: text mining and document analysis; image and video data analysis; data processing and data mining; recommendation algorithms and systems; probabilistic models and applications; knowledge engineering applications; and knowledge graph and knowledge management. The papers of the second volume are organized in the following topical sections: constraints and satisfiability; formal reasoning and ontologies; deep learning; network knowledge representation and learning; and social knowledge analysis and management.

This book seeks to promote the exploitation of data science in healthcare systems. The focus is on advancing the automated analytical methods used to extract new knowledge from data for healthcare applications. To do so, the book draws on several interrelated disciplines, including machine learning, big data analytics, statistics, pattern recognition, computer vision, and Semantic Web technologies, and focuses on their direct application to healthcare. Building on three tutorial-like chapters on data science in healthcare, the following eleven chapters highlight success stories on the application of data science in healthcare, where data science and artificial intelligence technologies have proven to be very promising. This book is primarily intended for data scientists involved in the healthcare or medical sector. By reading this book, they will gain essential insights into the modern data science technologies needed to advance innovation for both healthcare businesses and patients. A basic grasp of data science is recommended in order to fully benefit from this book.

The 30-volume set, comprising the LNCS books 12346 until 12375, constitutes the refereed proceedings of the 16th European Conference on Computer Vision, ECCV 2020, which was planned to be held in Glasgow, UK, during August 23-28, 2020. The conference was held virtually due to the COVID-19 pandemic. The 1360 revised papers presented in these proceedings were carefully reviewed and selected from a total of 5025 submissions. The papers deal with topics such as computer vision; machine learning; deep neural networks; reinforcement learning; object recognition; image classification; image processing; object detection; semantic segmentation; human pose estimation; 3d reconstruction; stereo vision; computational photography; neural networks; image coding; image reconstruction; object recognition; motion estimation.

This thesis contributes to the literature of understanding and recognizing human activities in videos. More specifically, the thesis draw line between short-range atomic actions and long-range complex activities. For the classification of the latter, the mainstream approach in literature is to divide the activity into a handful of short segments, called atomic actions. Then, a neural model, such as 3D CNN, is trained to represent and classify each segment independently. Then, the activity-level classification probability scores are obtained by pooling over that of the segments. Differently, this work argues that long-range activities are better classified in full. That is to say, the neural model has to reason about the long-range activity, all at once, to better recognize it. Based on this argument, the thesis proposes different methods and neural network models for recognizing these complex activities.

This book constitutes the refereed proceedings of the 9th International Conference on Bioinspired Optimization Methods and Their Applications, BIOMA 2020, held in Brussels, Belgium, in November 2020. The 24 full papers presented in this book were carefully reviewed and selected from 68 submissions. The papers in this BIOMA proceedings specialized in bioinspired algorithms as a means for solving the optimization problems and came in two categories: theoretical studies and methodology advancements on the one hand, and algorithm adjustments and their applications on the other. Due to the Corona pandemic BIOMA 2020 was held as a virtual event.

Many AI (and machine learning) tasks present in dual forms, e.g., English-to-Chinese translation vs. Chinese-to-English translation, speech

recognition vs. speech synthesis,question answering vs. question generation, and image classification vs. image generation. Dual learning is a new learning framework that leverages the primal-dual structure of AI tasks to obtain effective feedback or regularization signals in order to enhance the learning/inference process. Since it was first introduced four years ago, the concept has attracted considerable attention in multiple fields, and been proven effective in numerous applications, such as machine translation, image-to-image translation, speech synthesis and recognition, (visual) question answering and generation, image captioning and generation, and code summarization and generation. Offering a systematic and comprehensive overview of dual learning, this book enables interested researchers (both established and newcomers) and practitioners to gain a better understanding of the state of the art in the field. It also provides suggestions for further reading and tools to help readers advance the area. The book is divided into five parts. The first part gives a brief introduction to machine learning and deep learning. The second part introduces the algorithms based on the dual reconstruction principle using machine translation, image translation, speech processing and other NLP/CV tasks as the demo applications. It covers algorithms, such as dual semi-supervised learning, dual unsupervised learning and multi-agent dual learning. In the context of image translation, it introduces algorithms including CycleGAN, DualGAN, DiscoGAN cdGAN and more recent techniques/applications. The third part presents various work based on the probability principle, including dual supervised learning and dual inference based on the joint-probability principle and dual semi-supervised learning based on the marginal-probability principle. The fourth part reviews various theoretical studies on dual learning and discusses its connections to other learning paradigms. The fifth part provides a summary and suggests future research directions.

This book constitutes the refereed proceedings of the 10th International Symposium on Parallel Architectures, Algorithms and Programming, PAAP 2019, held in Guangzhou, China, in December 2019. The 39 revised full papers and 8 revised short papers presented were carefully reviewed and selected from 121 submissions. The papers deal with research results and development activities in all aspects of parallel architectures, algorithms and programming techniques.

The eight-volume set LNCS 12901, 12902, 12903, 12904, 12905, 12906, 12907, and 12908 constitutes the refereed proceedings of the 24th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI 2021, held in Strasbourg, France, in September/October 2021.* The 531 revised full papers presented were carefully reviewed and selected from 1630 submissions in a double-blind review process. The papers are organized in the following topical sections: Part I: image segmentation Part II: machine learning - self-supervised learning; machine learning - semi-supervised learning; and machine learning - weakly supervised learning Part III: machine learning - advances in machine learning theory; machine learning - attention models; machine learning - domain adaptation; machine learning - federated learning; machine learning - interpretability / explainability; and machine learning - uncertainty Part IV: image registration; image-guided interventions and surgery; surgical data science; surgical planning and simulation; surgical skill and work flow analysis; and surgical visualization and mixed, augmented and virtual reality Part V: computer aided diagnosis; integration of imaging with non-imaging biomarkers; and outcome/disease prediction Part VI: image reconstruction; clinical applications - cardiac; and clinical applications - vascular Part VII: clinical applications - abdomen; clinical applications - breast; clinical applications - dermatology; clinical applications - fetal imaging; clinical applications - lung; clinical applications - neuroimaging - brain development; clinical applications - neuroimaging - DWI and tractography; clinical applications - neuroimaging - functional brain networks; clinical applications - neuroimaging – others; and clinical applications - oncology Part VIII: clinical applications - ophthalmology; computational (integrative) pathology; modalities - microscopy; modalities - histopathology; and modalities - ultrasound *The conference was held virtually.

Recent advances in computational algorithms, along with the advent of whole slide imaging as a platform for embedding artificial intelligence (AI), are transforming pattern recognition and image interpretation for diagnosis and prognosis. Yet most pathologists have just a passing knowledge of data mining, machine learning, and AI, and little exposure to the vast potential of these powerful new tools for medicine in general and pathology in particular. In Artificial Intelligence and Deep Learning in Pathology, Dr. Stanley Cohen covers the nuts and bolts of all aspects of machine learning, up to and including AI, bringing familiarity and understanding to pathologists at all levels of experience. Focuses heavily on applications in medicine, especially pathology, making unfamiliar material accessible and avoiding complex mathematics whenever possible. Covers digital pathology as a platform for primary diagnosis and augmentation via deep learning, whole slide imaging for 2D and 3D analysis, and general principles of image analysis and deep learning. Discusses and explains recent accomplishments such as algorithms used to diagnose skin cancer from photographs, AI-based platforms developed to identify lesions of the retina, using computer vision to interpret electrocardiograms, identifying mitoses in cancer using learning algorithms vs. signal processing algorithms, and many more.

This two-volume set, LNCS 11641 and 11642, constitutes the thoroughly refereed proceedings of the Third International Joint Conference, APWeb-WAIM 2019, held in Chengdu, China, in August 2019. The 42 full papers presented together with 17 short papers, and 6 demonstration papers were carefully reviewed and selected from 180 submissions. The papers are organized around the following topics: Big Data Analytics; Data and Information Quality; Data Mining and Application; Graph Data and Social Networks; Information Extraction and Retrieval; Knowledge Graph; Machine Learning; Recommender Systems; Storage, Indexing and Physical Database Design; Spatial, Temporal and Multimedia Databases; Text Analysis and Mining; and Demo.

This two volume set of LNAI 11108 and LNAI 11109 constitutes the refereed proceedings of the 7th CCF Conference on Natural Language Processing and Chinese Computing, NLPCC 2018, held in Hohhot, China, in August 2018. The 55 full papers and 31 short papers presented were carefully reviewed and selected from 308 submissions. The papers of the first volume are organized in the following topics: conversational Bot/QA/IR; knowledge graph/IE; machine learning for NLP; machine translation; and NLP applications. The papers of the second volume are organized as follows: NLP for social network; NLP fundamentals; text mining; and short papers.

This book constitutes the thoroughly refereed proceedings of the 6th National Conference of Social Media Processing, SMP 2017, held in Beijing, China, in September 2017. The 28 revised full papers presented were carefully reviewed and selected from 140 submissions. The papers address issues such as: knowledge discovery for data; natural language processing; text mining and sentiment analysis; social network analysis and social computing.

The 6-volume set, comprising the LNCS books 12535 until 12540, constitutes the refereed proceedings of 28 out of the 45 workshops held at the 16th European Conference on Computer Vision, ECCV 2020. The conference was planned to take place in Glasgow, UK, during August 23-28, 2020, but changed to a virtual format due to the COVID-19 pandemic. The 249 full papers, 18 short papers, and 21 further contributions included in the workshop proceedings were carefully reviewed and selected from a total of 467 submissions. The papers deal with diverse computer vision topics. Part II focusses on commands for autonomous vehicles; computer vision for ART analysis; sign language recognition, translation and production; visual inductive priors for data-efficient deep learning; 3D poses in the wild challenge; map-based localization for autonomous driving; recovering 6D object pose; and shape recovery from partial textured 3D scans.

This four-volume set of LNCS 12821, LNCS 12822, LNCS 12823 and LNCS 12824, constitutes the refereed proceedings of the 16th International Conference on Document Analysis and Recognition, ICDAR 2021, held in Lausanne,

Switzerland in September 2021. The 182 full papers were carefully reviewed and selected from 340 submissions, and are presented with 13 competition reports. The papers are organized into the following topical sections: document analysis for literature search, document summarization and translation, multimedia document analysis, mobile text recognition, document analysis for social good, indexing and retrieval of documents, physical and logical layout analysis, recognition of tables and formulas, and natural language processing (NLP) for document understanding.

Intelligent Speech Signal Processing investigates the utilization of speech analytics across several systems and real-world activities, including sharing data analytics related information, creating collaboration networks between several participants, and implementing video-conferencing in different application areas. It provides a forum for readers to discover the characteristics of intelligent speech signal processing systems across different domains. Chapters focus on the latest applications of speech data analysis and management tools across different recording systems. The book emphasizes the multi-disciplinary nature of the field, presenting different applications and challenges with extensive studies on the design, implementation, development, and management of intelligent systems, neural networks, and related machine learning techniques for speech signal processing. Highlights different data analytics techniques in speech signal processing, including machine learning, and data mining Illustrates different applications and challenges across the design, implementation, and management of intelligent systems and neural networks techniques for speech signal processing Includes coverage of biomodal speech recognition, voice activity detection, spoken language and speech disorder identification, automatic speech to speech summarization, and convolutional neural networks

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.

This book constitutes the proceedings of the 17th China National Conference on Computational Linguistics, CCL 2018, and the 6th International Symposium on Natural Language Processing Based on Naturally Annotated Big Data, NLP-NABD 2018, held in Changsha, China, in October 2018. The 33 full papers presented in this volume were carefully reviewed and selected from 84 submissions. They are organized in topical sections named: Semantics; machine translation; knowledge graph and information extraction; linguistic resource annotation and evaluation; information retrieval and question answering; text classification and summarization; social computing and sentiment analysis; and NLP applications.

Over the past years, businesses have had to tackle the issues caused by numerous forces from political, technological and societal environment. The changes in the global market and increasing uncertainty require us to focus on disruptive innovations and to investigate this phenomenon from different perspectives. The benefits of innovations are related to lower costs, improved efficiency, reduced risk, and better response to the customers' needs due to new products, services or processes. On the other hand, new business models expose various risks, such as cyber risks, operational risks, regulatory risks, and others. Therefore, we believe that the entrepreneurial behavior and global mindset of decision-makers significantly contribute to the development of innovations, which benefit by closing the prevailing gap between developed and developing countries. Thus, this Special Issue contributes to closing the research gap in the literature by providing a platform for a scientific debate on innovation, internationalization and entrepreneurship, which would facilitate improving the resilience of businesses to future disruptions. Order Your Print Copy