a neural probabilistic language model

A Neural Probabilistic Language Model: A Deep Dive into the Architecture and Applications

Author: Dr. Emily Carter, PhD, Professor of Computer Science and Linguistics, Stanford University. Dr. Carter has over 15 years of experience in natural language processing (NLP), with a particular focus on the development and application of neural probabilistic language models. Her research has been published in leading journals and conferences, and she is a recipient of the prestigious ACM SIGIR Best Paper Award.

Publisher: MIT Press, a renowned publisher with a long-standing reputation for excellence in computer science and artificial intelligence publications. Their rigorous peer-review process ensures the quality and accuracy of the published work, making them a credible source for information on advanced topics such as a neural probabilistic language model.

Editor: Dr. David Chiang, PhD, Professor of Computer Science, University of California, Berkeley. Dr. Chiang's expertise lies in computational linguistics and machine translation, with significant contributions to the field of statistical and neural language modeling. His editorial oversight guarantees the technical accuracy and relevance of the article.

Abstract: This report provides a comprehensive overview of a neural probabilistic language model (NPLM), exploring its architecture, training methodologies, and diverse applications. We delve into the advancements made in this field, examining key research findings and comparing different NPLM architectures. The report also addresses the challenges and limitations of NPLMs, and discusses potential avenues for future research.

1. Introduction: The Rise of Neural Probabilistic Language Models

Language modeling is the fundamental task of assigning probabilities to sequences of words. Traditional statistical language models, like n-grams, relied on counting word occurrences in large corpora. However, these methods struggle with long-range dependencies and the sparsity of data. The advent of neural networks revolutionized the field, leading to the development of a neural probabilistic language model. These models leverage the power of neural networks to capture complex relationships between words and learn rich representations of language. A neural probabilistic language model significantly outperforms traditional methods in various downstream NLP tasks.

1. Architecture of a Neural Probabilistic Language Model

A core component of a neural probabilistic language model is its ability to represent words as dense, low-dimensional vectors, often called word embeddings. These embeddings capture semantic and syntactic information about words, allowing the model to understand relationships between them. Popular architectures include:

Recurrent Neural Networks (RNNs): RNNs, particularly Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs), are well-suited for processing sequential data like text. They maintain a hidden state that is updated at each time step, allowing them to capture long-range dependencies. A neural probabilistic language model based on RNNs excels in tasks requiring contextual understanding.

Transformers: Transformers, introduced in the seminal paper "Attention is All You Need," have achieved state-of-the-art results in various NLP tasks. They rely on the attention mechanism, which allows the model to weigh the importance of different words in the input sequence when predicting the next word. This architecture is particularly effective in capturing long-range dependencies and parallel processing. A neural probabilistic language model using a transformer architecture often exhibits superior performance.

Convolutional Neural Networks (CNNs): CNNs can also be used to build a neural probabilistic language model, effectively capturing local n-gram information. While less commonly used than RNNs and transformers, they offer advantages in terms of computational efficiency for certain tasks.

1. Training a Neural Probabilistic Language Model

Training a neural probabilistic language model involves optimizing the model's parameters to maximize the likelihood of a large text corpus. This is typically done using stochastic gradient descent (SGD) or its variants, such as Adam. The training process requires a significant amount of computational resources and time, especially for large models and datasets. Regularization techniques, such as dropout and weight decay, are often employed to prevent overfitting.

1. Evaluating a Neural Probabilistic Language Model

The performance of a neural probabilistic language model is typically evaluated using metrics like perplexity. Perplexity measures the uncertainty of the model in predicting the next word in a sequence.

Lower perplexity indicates better performance. Other evaluation metrics include BLEU score (for machine translation) and ROUGE score (for text summarization), which are task-specific metrics.

1. Applications of a Neural Probabilistic Language Model

A neural probabilistic language model has a wide range of applications across various NLP tasks:

Machine Translation: A neural probabilistic language model is fundamental to modern machine translation systems, enabling accurate and fluent translation between languages.

Text Summarization: A neural probabilistic language model can generate concise and informative summaries of longer texts.

Question Answering: A neural probabilistic language model can be used to build systems that answer questions posed in natural language.

Chatbots and Conversational AI: A neural probabilistic language model is crucial for building engaging and coherent chatbots.

Sentiment Analysis: A neural probabilistic language model can be used to determine the sentiment expressed in a piece of text.

Speech Recognition: A neural probabilistic language model plays a vital role in improving the accuracy of speech recognition systems.

1. Challenges and Limitations

Despite their impressive performance, a neural probabilistic language model faces several challenges:

Computational Cost: Training large NPLMs requires significant computational resources and time.

Data Requirements: NPLMs require massive amounts of training data to achieve good performance.

Bias and Fairness: NPLMs can inherit biases present in the training data, leading to unfair or discriminatory outcomes.

Interpretability: Understanding the internal workings of a complex NPLM can be challenging.

1. Future Directions

Future research on a neural probabilistic language model will likely focus on:

Improving Efficiency: Developing more efficient training algorithms and architectures.

Addressing Bias: Developing techniques to mitigate bias in NPLMs.

Enhancing Interpretability: Developing methods to make NPLMs more interpretable.

Transfer Learning: Exploring the use of pre-trained NPLMs for various downstream tasks.

1. Conclusion

A neural probabilistic language model represents a significant advancement in the field of natural language processing. Its ability to learn rich representations of language and capture complex relationships between words has led to remarkable improvements in various NLP tasks. While challenges remain, ongoing research continues to push the boundaries of NPLM capabilities, promising even more impressive advancements in the years to come. The development of more efficient, robust, and interpretable NPLMs is crucial for unlocking the full potential of language technology.

FAQs:

1. What is the difference between a traditional n-gram model and a neural probabilistic language model? N-gram models rely on simple word counts, limiting their ability to capture long-range dependencies and semantic relationships. Neural probabilistic language models use neural networks to learn rich representations of words and capture complex contextual information.

1. What are word embeddings, and why are they important in NPLMs? Word embeddings are dense

vector representations of words that capture semantic and syntactic information. They allow the model to understand relationships between words and improve performance on various NLP tasks.

1. What is perplexity, and how is it used to evaluate NPLMs? Perplexity is a metric that measures the uncertainty of a language model in predicting the next word in a sequence. Lower perplexity indicates better performance.

1. What are the main architectures used in NPLMs? Common architectures include Recurrent Neural Networks (RNNs, LSTMs, GRUs), Transformers, and Convolutional Neural Networks (CNNs).

1. What are some common challenges in training NPLMs? Challenges include high computational cost, the need for large datasets, potential biases in training data, and difficulties in interpreting model decisions.

1. How can biases in NPLMs be mitigated? Techniques to mitigate bias include careful data curation, bias detection algorithms, and adversarial training methods.

1. What are some applications of NPLMs beyond machine translation? NPLMs are used in text summarization, question answering, chatbots, sentiment analysis, and speech recognition.

1. What is the role of attention mechanisms in transformer-based NPLMs? Attention mechanisms allow the model to focus on different parts of the input sequence when making predictions, improving its ability to capture long-range dependencies.

1. What are some future research directions in NPLM development? Future research focuses on efficiency improvements, bias mitigation, enhanced interpretability, and transfer learning

applications.

Related Articles:

1. "Attention is All You Need" - This seminal paper introduced the transformer architecture, revolutionizing the field of NPLMs.

1. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding" - This paper introduced BERT, a powerful pre-trained NPLM that has achieved state-of-the-art results on many NLP tasks.

1. "A Neural Probabilistic Language Model" by Bengio et al. - A foundational paper in the development of NPLMs.

1. "Recurrent Neural Network Architectures for Language Modeling" - Explores the use of RNNs in NPLMs.

1. "Convolutional Neural Networks for Sentence Classification" - Discusses the use of CNNs in language modeling.

1. "Evaluating the Performance of Neural Language Models" - Examines various evaluation metrics for NPLMs.

1. "Addressing Bias in Neural Language Models" - Focuses on techniques to mitigate bias in NPLMs.

1. "Improving the Efficiency of Neural Language Models" - Explores strategies for improving the computational efficiency of NPLMs.

1. "Transfer Learning for Natural Language Processing" - Discusses the application of pre-trained NPLMs to various downstream tasks.

a neural probabilistic language model: Neural Representations of Natural Language Lyndon White, Roberto Togneri, Wei Liu, Mohammed Bennamoun, 2018-08-29 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 a 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.

a neural probabilistic language model: Handbook of Probabilistic Models Pijush Samui, Dieu Tien Bui, Subrata Chakraborty, Ravinesh Deo, 2019-10-05 Handbook of Probabilistic Models carefully examines the application of advanced probabilistic models in conventional engineering fields. In this comprehensive handbook, practitioners, researchers and scientists will find detailed explanations of technical concepts, applications of the proposed methods, and the respective scientific approaches needed to solve the problem. This book provides an interdisciplinary approach that creates advanced probabilistic models for engineering fields, ranging from conventional fields of mechanical engineering and civil engineering, to electronics, electrical, earth sciences, climate, agriculture, water resource, mathematical sciences and computer sciences. Specific topics covered include minimax probability machine regression, stochastic finite element method, relevance vector machine, logistic regression, Monte Carlo simulations, random matrix, Gaussian process regression, Kalman filter, stochastic optimization, maximum likelihood, Bayesian inference, Bayesian update, kriging, copula-statistical models, and more. - Explains the application of advanced probabilistic

models encompassing multidisciplinary research - Applies probabilistic modeling to emerging areas in engineering - Provides an interdisciplinary approach to probabilistic models and their applications, thus solving a wide range of practical problems

a neural probabilistic language model: Neural Networks: Tricks of the Trade Grégoire Montavon, Geneviève Orr, Klaus-Robert Müller, 2012-11-14 The twenty last years have been marked by an increase in available data and computing power. In parallel to this trend, the focus of neural network research and the practice of training neural networks has undergone a number of important changes, for example, use of deep learning machines. The second edition of the book augments the first edition with more tricks, which have resulted from 14 years of theory and experimentation by some of the world's most prominent neural network researchers. These tricks can make a substantial difference (in terms of speed, ease of implementation, and accuracy) when it comes to putting algorithms to work on real problems.

a neural probabilistic language model: Probabilistic Machine Learning Kevin P. Murphy, 2022-03-01 A detailed and up-to-date introduction to machine learning, presented through the unifying lens of probabilistic modeling and Bayesian decision theory. This book offers a detailed and up-to-date introduction to machine learning (including deep learning) through the unifying lens of probabilistic modeling and Bayesian decision theory. The book covers mathematical background (including linear algebra and optimization), basic supervised learning (including linear and logistic regression and deep neural networks), as well as more advanced topics (including transfer learning and unsupervised learning). End-of-chapter exercises allow students to apply what they have learned, and an appendix covers notation. Probabilistic Machine Learning grew out of the author’s 2012 book, Machine Learning: A Probabilistic Perspective. More than just a simple update, this is a completely new book that reflects the dramatic developments in the field since 2012, most notably deep learning. In addition, the new book is accompanied by online Python code, using libraries such as scikit-learn, JAX, PyTorch, and Tensorflow, which can be used to reproduce nearly all the figures; this code can be run inside a web browser using cloud-based notebooks, and provides a practical complement to the theoretical topics discussed in the book. This introductory text will be followed by a sequel that covers more advanced topics, taking the same probabilistic approach.

a neural probabilistic language model: Neural Machine Translation Philipp Koehn, 2020-06-18 Learn how to build machine translation systems with deep learning from the ground up, from basic concepts to cutting-edge research.

a neural probabilistic language model: Deep Learning Ian Goodfellow, Yoshua Bengio, Aaron Courville, 2016-11-10 An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by

undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

a neural probabilistic language model: Text, Speech, and Dialogue Ivan Habernal, Vaclav Matousek, 2013-08-17 This book constitutes the refereed proceedings of the 16th International Conference on Text, Speech and Dialogue, TSD 2013, held in Pilsen, Czech Republic, in September 2013. The 65 papers presented together with 5 invited talks were carefully reviewed and selected from 148 submissions. The main topics of this year's conference was corpora, texts and transcription, speech analysis, recognition and synthesis, and their intertwining within NL dialogue systems. The topics also included speech recognition, corpora and language resources, speech and spoken language generation, tagging, classification and parsing of text and speech, semantic processing of text and speech, integrating applications of text and speech processing, as well as automatic dialogue systems, and multimodal techniques and modelling.

a neural probabilistic language model: *Introduction to Python and Large Language Models* Dilyan Grigorov,

a neural probabilistic language model: Neural Network Methods for Natural Language Processing Yoav Goldberg, 2022-06-01 Neural networks are a family of powerful machine learning models. This book focuses on the application of neural network models 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.

a neural probabilistic language model: Neural Information Processing Long Cheng, Andrew Chi Sing Leung, Seiichi Ozawa, 2018-12-03 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 third volume, LNCS 11303, is organized in topical sections on embedded learning, transfer learning, reinforcement learning, and other learning approaches.

a neural probabilistic language model: Representation Learning for Natural Language Processing Zhiyuan Liu, Yankai Lin, Maosong Sun, 2020-07-03 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.

a neural probabilistic language model: *Deep Learning for Natural Language Processing* Jason Brownlee, 2017-11-21 Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

a neural probabilistic language model: Speech & Language Processing Dan Jurafsky, 2000-09

a neural probabilistic language model: Quick Start Guide to Large Language Models Sinan Ozdemir, 2024-09-26 The Practical, Step-by-Step Guide to Using LLMs at Scale in Projects and Products Large Language Models (LLMs) like Llama 3, Claude 3, and the GPT family are demonstrating breathtaking capabilities, but their size and complexity have deterred many practitioners from applying them. In Quick Start Guide to Large Language Models, Second Edition, pioneering data scientist and AI entrepreneur Sinan Ozdemir clears away those obstacles and provides a guide to working with, integrating, and deploying LLMs to solve practical problems. Ozdemir brings together all you need to get started, even if you have no direct experience with LLMs: step-by-step instructions, best practices, real-world case studies, and hands-on exercises. Along the way, he shares insights into LLMs' inner workings to help you optimize model choice, data formats, prompting, fine-tuning, performance, and much more. The resources on the companion website include sample datasets and up-to-date code for working with open- and closed-source LLMs such as those from OpenAI (GPT-4 and GPT-3.5), Google (BERT, T5, and Gemini), X (Grok), Anthropic (the Claude family), Cohere (the Command family), and Meta (BART and the LLaMA family). Learn key concepts: pre-training, transfer learning, fine-tuning, attention, embeddings, tokenization, and more Use APIs and Python to fine-tune and customize LLMs for your requirements Build a complete neural/semantic information retrieval system and attach to conversational LLMs for building retrieval-augmented generation (RAG) chatbots and AI Agents Master advanced prompt engineering techniques like output structuring, chain-of-thought prompting, and semantic few-shot prompting Customize LLM embeddings to build a complete recommendation engine from scratch with user data that outperforms out-of-the-box embeddings from OpenAI Construct and fine-tune multimodal Transformer architectures from scratch using open-source LLMs and large visual datasets Align LLMs using Reinforcement Learning from Human and AI Feedback (RLHF/RLAIF) to build conversational agents from open models like Llama 3 and FLAN-T5 Deploy prompts and custom fine-tuned LLMs to the cloud with scalability and evaluation pipelines in mind Diagnose and optimize LLMs for speed, memory, and performance with quantization, probing, benchmarking, and evaluation frameworks A refreshing and inspiring resource. Jam-packed with practical guidance and clear explanations that leave you smarter about this incredible new field. --Pete Huang, author of The Neuron Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

a neural probabilistic language model: Web Information Systems and Applications Guojun Wang, Xuemin Lin, James Hendler, Wei Song, Zhuoming Xu, Genggeng Liu, 2020-09-22 This book constitutes the proceedings of the 17th International Conference on Web Information Systems and Applications, WISA 2020, held in Guangzhou, China, in September 2020. The 42 full papers and 16 short papers presented were carefully reviewed and selected from 165 submissions. The papers are grouped in topical sections on world wide web, recommendation, query processing and algorithm, natural language processing, machine learning, graph query, edge computing and data mining, data privacy and security, and blockchain.

a neural probabilistic language model: *Web Information Systems and Applications* Cheqing Jin,

a neural probabilistic language model: *Python Deep Learning* Valentino Zocca, Gianmario

Spacagna, Daniel Slater, Peter Roelants, 2017-04-28 Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book Is For This book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H20. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

a neural probabilistic language model: Deep Learning in Natural Language Processing Li Deng, Yang Liu, 2018-05-23 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.

a neural probabilistic language model: Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data Maosong Sun, Zhiyuan Liu, Min Zhang, Yang Liu, 2015-11-07 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.

a neural probabilistic language model: Turkish Natural Language Processing Kemal Oflazer, Murat Saraçlar, 2018-07-20 This book brings together work on Turkish natural language and speech processing over the last 25 years, covering numerous fundamental tasks ranging from morphological processing and language modeling, to full-fledged deep parsing and machine translation, as well as computational resources developed along the way to enable most of this work. Owing to its complex morphology and free constituent order, Turkish has proved to be a fascinating language for natural language and speech processing research and applications. After an overview of the aspects of Turkish that make it challenging for natural language and speech processing tasks, this book discusses in detail the main tasks and applications of Turkish natural language and speech processing. A compendium of the work on Turkish natural language and speech processing, it is a valuable reference for new researchers considering computational work on Turkish, as well as a one-stop resource for commercial and research institutions planning to develop applications for Turkish. It also serves as a blueprint for similar work on other Turkic languages such as Azeri, Turkmen and Uzbek.

a neural probabilistic language model: Intelligent Decision Technologies Ireneusz Czarnowski, Robert J. Howlett, Lakhmi C. Jain, 2021-07-07 This book contains selected papers from the KES-IDT-2021 conference, being held as a virtual conference in June 14–16, 2021. The KES-IDT is an interdisciplinary conference with opportunities for the presentation of new research results and discussion about them under the common title Intelligent Decision Technologies. The conference has been creating for years a platform for knowledge transfer and the generation of new ideas in the field of intelligent decision making. The range of topics discussed during the conference covered methods of classification, prediction, data analysis, big data, decision support, knowledge engineering, modeling, social networks and many more in areas such as finance, economy, management and transportation. The discussed topics covered also decision making for problems regarding the electric vehicle industry. The book contains also several sections devoted to specific topics, such as Advances in intelligent data processing and its applications Multi-criteria decision analysis methods Knowledge engineering in large-scale systems High-dimensional data analysis Spatial data analysis and sparse estimation Innovative technologies and applications in computer intelligence Intelligent diagnosis and monitoring of systems Decision making theory for economics.

a neural probabilistic language model: *Social Computing* Wanxiang Che, Qilong Han, Hongzhi Wang, Weipeng Jing, Shaoliang Peng, Junyu Lin, Guanglu Sun, Xianhua Song, Hongtao Song, Zeguang Lu, 2016-07-30 This two volume set (CCIS 623 and 634) constitutes the refereed proceedings of the Second International Conference of Young Computer Scientists, Engineers and Educators, ICYCSEE 2016, held in Harbin, China, in August 2016. The 91 revised full papers presented were carefully reviewed and selected from 338 submissions. The papers are organized in topical sections on Research Track (Part I) and Education Track, Industry Track, and Demo Track (Part II) and cover a wide range of topics related to social computing, social media, social network analysis, social modeling, social recommendation, machine learning, data mining.

a neural probabilistic language model: Computer and Information Science Roger Lee, 2019-08-06 This book gathers the outcomes of the 18th IEEE/ACIS International Conference on Computer and Information Science (ICIS 2019), which was held on June 12–14, 2019 in Beijing, China. The aim of the conference was to bring together researchers and scientists, businesspeople and entrepreneurs, teachers, engineers, computer users, and students to discuss the various fields of computer science and to share their experiences and exchange new ideas and information in a meaningful way. Further, they presented research results on all aspects (theory, applications and tools) of computer and information science, and discussed the practical challenges encountered in their work and the solutions they adopted to overcome them. The book highlights the best papers from those accepted for presentation at the conference. They were chosen based on review scores submitted by members of the program committee and underwent further rigorous rounds of review.

From this second round, 15 of the conference's most promising papers were selected for this Springer (SCI) book and not the conference proceedings. We eagerly await the important contributions that we know these authors will make to the field of computer and information science.

a neural probabilistic language model: Legal Knowledge and Information Systems M. Palmirani, 2018-12-12 Artificial intelligence as applied to the legal domain has gained momentum thanks to the large, annotated corporate legal and case-law collections, human chats, and social media information now available in open data. Often represented in XML or other Semantic Web technologies, these now make it possible to use the AI theory developed by the JURIX community in over thirty years of research. Innovative machine and deep-learning techniques with which to classify legal texts and detect terms, principles, concepts, evidence, named entities, and rules are also emerging, and the last five years have seen a gradual increase in their practical application. This book presents papers from the 31st International Conference on Legal Knowledge and Information Systems (JURIX 2018), held in Groningen, the Netherlands, in December 2018. The support of the Dutch Foundation for Legal Knowledge Based Systems for the JURIX conference has transformed a domestic workshop into an international event, with theoretical contributions, applied work, demo prototypes, a hackathon, and a doctoral consortium. Of the 72 submissions received, 17 full papers and 11 short papers were selected for publication, representing an acceptance rate of approximately 38%. Machine learning for the legal domain prevails in the JURIX 2018 program, with traditional research mainstreams concerning legal reasoning and argumentation, natural-language processing, legal-text retrieval, and legal semantic modelling. An emerging topic is blockchain, which has graduated from the workshop area to the main program. The book offers an overview of the ways in which innovative information technologies are merging with legal theory, argumentation, and practice.

a neural probabilistic language model: PRICAI 2014: Trends in Artificial Intelligence Duc-Nghia Pham, Seong-Bae Park, 2014-11-12 This book constitutes the refereed proceedings of the 13th Pacific Rim Conference on Artificial Intelligence, PRICAI 2014, held in Gold Coast, Queensland, Australia, in December 2014. The 74 full papers and 20 short papers presented in this volume were carefully reviewed and selected from 203 submissions. The topics include inference; reasoning; robotics; social intelligence. AI foundations; applications of AI; agents; Bayesian networks; neural networks; Markov networks; bioinformatics; cognitive systems; constraint satisfaction; data mining and knowledge discovery; decision theory; evolutionary computation; games and interactive entertainment; heuristics; knowledge acquisition and ontology; knowledge representation, machine learning; multimodal interaction; natural language processing; planning and scheduling; probabilistic.

a neural probabilistic language model: Innovations in Machine Learning Dawn E. Holmes, 2006-03-09 Machine learning is currently one of the most rapidly growing areas of research in computer science. In compiling this volume we have brought together contributions from some of the most prestigious researchers in this field. This book covers the three main learning systems; symbolic learning, neural networks and genetic algorithms as well as providing a tutorial on learning casual influences. Each of the nine chapters is self-contained. Both theoreticians and application scientists/engineers in the broad area of artificial intelligence will find this volume valuable. It also provides a useful sourcebook for Postgraduate since it shows the direction of current research.

a neural probabilistic language model: Deep Learning Ian Goodfellow, Yoshua Bengio, Aaron Courville, 2016-11-18 An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn

complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

a neural probabilistic language model: Information Theory and Language Łukasz Dębowski , Christian Bentz, 2020-12-15 “Information Theory and Language” is a collection of 12 articles that appeared recently in Entropy as part of a Special Issue of the same title. These contributions represent state-of-the-art interdisciplinary research at the interface of information theory and language studies. They concern in particular: • Applications of information theoretic concepts such as Shannon and Rényi entropies, mutual information, and rate–distortion curves to the research of natural languages; • Mathematical work in information theory inspired by natural language phenomena, such as deriving moments of subword complexity or proving continuity of mutual information; • Empirical and theoretical investigation of quantitative laws of natural language such as Zipf’s law, Herdan’s law, and Menzerath–Altmann’s law; • Empirical and theoretical investigations of statistical language models, including recently developed neural language models, their entropies, and other parameters; • Standardizing language resources for statistical investigation of natural language; • Other topics concerning semantics, syntax, and critical phenomena. Whereas the traditional divide between probabilistic and formal approaches to human language, cultivated in the disjoint scholarships of natural sciences and humanities, has been blurred in recent years, this book can contribute to pointing out potential areas of future research cross-fertilization.

a neural probabilistic language model: Deep Learning: Fundamentals, Theory and Applications Kaizhu Huang, Amir Hussain, Qiu-Feng Wang, Rui Zhang, 2019-02-15 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.

a neural probabilistic language model: *Making Presentation Math Computable* André Greiner-Petter, 2023-01-24 This Open-Access-book addresses the issue of translating mathematical

expressions from LaTeX to the syntax of Computer Algebra Systems (CAS). Over the past decades, especially in the domain of Sciences, Technology, Engineering, and Mathematics (STEM), LaTeX has become the de-facto standard to typeset mathematical formulae in publications. Since scientists are generally required to publish their work, LaTeX has become an integral part of today's publishing workflow. On the other hand, modern research increasingly relies on CAS to simplify, manipulate, compute, and visualize mathematics. However, existing LaTeX import functions in CAS are limited to simple arithmetic expressions and are, therefore, insufficient for most use cases. Consequently, the workflow of experimenting and publishing in the Sciences often includes time-consuming and error-prone manual conversions between presentational LaTeX and computational CAS formats. To address the lack of a reliable and comprehensive translation tool between LaTeX and CAS, this thesis makes the following three contributions. First, it provides an approach to semantically enhance LaTeX expressions with sufficient semantic information for translations into CAS syntaxes. Second, it demonstrates the first context-aware LaTeX to CAS translation framework LaCASt. Third, the thesis provides a novel approach to evaluate the performance for LaTeX to CAS translations on large-scaled datasets with an automatic verification of equations in digital mathematical libraries. This is an open access book.

a neural probabilistic language model: Statistical Foundations of Actuarial Learning and its Applications Mario V. Wüthrich, Michael Merz, 2022-11-22 This open access book discusses the statistical modeling of insurance problems, a process which comprises data collection, data analysis and statistical model building to forecast insured events that may happen in the future. It presents the mathematical foundations behind these fundamental statistical concepts and how they can be applied in daily actuarial practice. Statistical modeling has a wide range of applications, and, depending on the application, the theoretical aspects may be weighted differently: here the main focus is on prediction rather than explanation. Starting with a presentation of state-of-the-art actuarial models, such as generalized linear models, the book then dives into modern machine learning tools such as neural networks and text recognition to improve predictive modeling with complex features. Providing practitioners with detailed guidance on how to apply machine learning methods to real-world data sets, and how to interpret the results without losing sight of the mathematical assumptions on which these methods are based, the book can serve as a modern basis for an actuarial education syllabus.

a neural probabilistic language model: Natural Language Processing and Chinese Computing Guodong Zhou, Juanzi Li, Dongyan Zhao, Yansong Feng, 2013-10-01 This book constitutes the refereed proceedings of the Second CCF Conference on Natural Language Processing and Chinese Computing, NLPCC 2013, held in Chongqing, China, during November 2013. The 31 revised full papers presented together with three keynote talks and 13 short papers were carefully reviewed and selected from 203 submissions. The papers are organized in topical sections on fundamentals on language computing; applications on language computing; machine learning for NLP; machine translation and multi-lingual information access; NLP for social media and web mining, knowledge acquisition; NLP for search technology and ads; NLP fundamentals; NLP applications; NLP for social media.

a neural probabilistic language model: Web and Big Data Leong Hou U, Haoran Xie, 2018-10-20 This book constitutes the thoroughly refereed post-conference proceedings of the First APWeb-WAIM 2018 Workshops, held jointly with the Second International Joint Conference APWeb-WAIM 2018 in Macau, China, in July 2018. The 31 full papers presented were carefully reviewed and selected from 44 submissions. The papers originating from five workshops present cutting-edge ideas, results, experiences, techniques, and tools from all aspects of web data management with the focus on mobile web data analytics; knowledge graph management and analysis; data management and mining on MOOCs; Big data analytics for healthcare; data science.

a neural probabilistic language model: Natural Language Processing and Chinese Computing Min Zhang, Vincent Ng, Dongyan Zhao, Sujian Li, Hongying Zan, 2018-08-13 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.

a neural probabilistic language model: *Computational Knowledge Vision* Wenbo Zheng, Fei-Yue Wang, 2024-08-19 Computational Knowledge Vision: The First Footprints presents a novel, advanced framework which combines structuralized knowledge and visual models. In advanced image and visual perception studies, a visual model's understanding and reasoning ability often determines whether it works well in complex scenarios. This book presents state-of-the-art mainstream vision models for visual perception. As computer vision is one of the key gateways to artificial intelligence and a significant component of modern intelligent systems, this book delves into computer vision systems that are highly specialized and very limited in their ability to do visual reasoning and causal inference. Questions naturally arise in this arena, including (1) How can human knowledge be incorporated with visual models? (2) How does human knowledge promote the performance of visual models? To address these problems, this book proposes a new framework for computer vision–computational knowledge vision. - Presents a concept and basic framework of Computational Knowledge Vision that extends the knowledge engineering methodology to the computer vision field - Discusses neural networks, meta-learning, graphs, and Transformer models - Illustrates a basic framework for Computational Knowledge Vision whose essential techniques include structuralized knowledge, knowledge projection, and conditional feedback

a neural probabilistic language model: Advanced Deep Learning with Python Ivan Vasilev, 2019-12-12 Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key FeaturesGet to grips with building faster and more robust deep learning architecturesInvestigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorchApply deep neural networks (DNNs) to computer vision problems, NLP, and GANsBook Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learnCover advanced and state-of-the-art neural network architecturesUnderstand the theory and math behind neural networksTrain DNNs and apply them to modern deep learning problemsUse CNNs for object detection and image segmentationImplement generative adversarial networks (GANs) and variational autoencoders to generate new imagesSolve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence modelsUnderstand DL techniques, such as meta-learning and graph neural networksWho this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with

advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

a neural probabilistic language model: Bayesian Speech and Language Processing Shinji Watanabe, Jen-Tzung Chien, 2015-07-15 A practical and comprehensive guide on how to apply Bayesian machine learning techniques to solve speech and language processing problems.

a neural probabilistic language model: Multilingual Natural Language Processing Applications Daniel Bikel, Imed Zitouni, 2012-05-11 Multilingual Natural Language Processing Applications is the first comprehensive single-source guide to building robust and accurate multilingual NLP systems. Edited by two leading experts, it integrates cutting-edge advances with practical solutions drawn from extensive field experience. Part I introduces the core concepts and theoretical foundations of modern multilingual natural language processing, presenting today's best practices for understanding word and document structure, analyzing syntax, modeling language, recognizing entailment, and detecting redundancy. Part II thoroughly addresses the practical considerations associated with building real-world applications, including information extraction, machine translation, information retrieval/search, summarization, question answering, distillation, processing pipelines, and more. This book contains important new contributions from leading researchers at IBM, Google, Microsoft, Thomson Reuters, BBN, CMU, University of Edinburgh, University of Washington, University of North Texas, and others. Coverage includes Core NLP problems, and today's best algorithms for attacking them Processing the diverse morphologies present in the world's languages Uncovering syntactical structure, parsing semantics, using semantic role labeling, and scoring grammaticality Recognizing inferences, subjectivity, and opinion polarity Managing key algorithmic and design tradeoffs in real-world applications Extracting information via mention detection, coreference resolution, and events Building large-scale systems for machine translation, information retrieval, and summarization Answering complex questions through distillation and other advanced techniques Creating dialog systems that leverage advances in speech recognition, synthesis, and dialog management Constructing common infrastructure for multiple multilingual text processing applications This book will be invaluable for all engineers, software developers, researchers, and graduate students who want to process large quantities of text in multiple languages, in any environment: government, corporate, or academic.

a neural probabilistic language model: Computational Data and Social Networks Andrea Tagarelli, Hanghang Tong, 2019-11-11 This book constitutes the refereed proceedings of the 8th International Conference on Computational Data and Social Networks, CSoNet 2019, held in Ho Chi Minh City, Vietnam, in November 2019. The 22 full and 8 short papers presented in this book were carefully reviewed and selected from 120 submissions. The papers appear under the following topical headings: Combinatorial Optimization and Learning; Influence Modeling, Propagation, and Maximization; NLP and Affective Computing; Computational Methods for Social Good; and User Profiling and Behavior Modeling.

a neural probabilistic language model: Natural Language Processing and Chinese Computing Xuanjing Huang, Jing Jiang, Dongyan Zhao, Yansong Feng, Yu Hong, 2018-01-03 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.

Additional Resources

NEURAL Definition & Meaning - Merriam-Webster
The meaning of NEURAL is of, relating to, or affecting a nerve or the nervous system. How to use neural in a sentence.

Neural DSP - Algorithmically Perfect
Open a new world of possibilities by loading Neural DSP plugins on your Quad Cortex

NEURAL | English meaning - Cambridge Dictionary
NEURAL definition: 1. involving a nerve or the system of nerves that includes the brain: 2. involving a nerve or the.... Learn more.

Neural network - Wikipedia
Neural networks are used to solve problems in artificial intelligence, and have thereby found applications in many disciplines, including predictive modeling, adaptive control, facial ...

What is a Neural Network? - GeeksforGeeks
Apr 3, 2025 · Neural networks are machine learning models that mimic the complex functions of the human brain. These models consist of interconnected nodes or neurons that process data, ...

neural, adj. & n. meanings, etymology and more | Oxford ...
What does the word neural mean? There are six meanings listed in OED's entry for the word neural. See ‘Meaning & use’ for definitions, usage, and quotation evidence. How common is ...

neural - Wiktionary, the free dictionary
May 27, 2025 · neural (comparative lebih neural, superlative paling neural) neural of, or relating to the nerves, neurons or the nervous system modelled on the arrangement of neurons in the brain

NEURAL Definition & Meaning - Merriam-Webster
The meaning of NEURAL is of, relating to, or affecting a nerve or the nervous system. How to use neural in a sentence.

Neural DSP - Algorithmically Perfect
Open a new world of possibilities by loading Neural DSP plugins on your Quad Cortex

NEURAL | English meaning - Cambridge Dictionary
NEURAL definition: 1. involving a nerve or the system of nerves that includes the brain: 2. involving a nerve or the.... Learn more.

Neural network - Wikipedia
Neural networks are used to solve problems in artificial intelligence, and have thereby found applications in many disciplines, including predictive modeling, adaptive control, facial ...

What is a Neural Network? - GeeksforGeeks
Apr 3, 2025 · Neural networks are machine learning models that mimic the complex functions of the human brain. These models consist of interconnected nodes or neurons that process data, ...

neural, adj. & n. meanings, etymology and more | Oxford ...
What does the word neural mean? There are six meanings listed in OED's entry for the word neural. See ‘Meaning & use’ for definitions, usage, and quotation evidence. How common is the ...

neural - Wiktionary, the free dictionary
May 27, 2025 · neural (comparative lebih neural, superlative paling neural) neural of, or relating to the nerves, neurons or the nervous system modelled on the arrangement of neurons in the brain

A Neural Probabilistic Language Model

A Neural Probabilistic Language Model

Related Articles

- a labeled diagram of the skeletal system
- a ladys guide to improper behavior
- a one piece game trading discord

Back to Home