

Neural Network Methods For Natural Language Processing (Synthesis Lectures On Human Language Technologies)

A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field’s linguistic and computational foundations. It is suitable for use in advanced undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

This undergraduate textbook introduces essential machine learning concepts in NLP in a unified and gentle mathematical framework.

Build your Machine Learning portfolio by creating 6 cutting-edge Artificial Intelligence projects using neural networks in Python Key FeaturesDiscover neural network architectures (like CNN and LSTM) that are driving recent advancements in AIBuild expert neural networks in Python using popular libraries such as KerasIncludes projects such as object detection, face identification, sentiment analysis, and moreBook Description Neural networks are at the core of recent AI advances, providing some of the best resolutions to many real-world problems, including image recognition, medical diagnosis, text analysis, and more. This book goes through some basic neural network and deep learning concepts, as well as some popular libraries in Python for implementing them. It contains practical demonstrations of neural networks in domains such as fare prediction, image classification, sentiment analysis, and more. In each case, the book provides a problem statement, the specific neural network architecture required to tackle that problem, the reasoning behind the algorithm used, and the associated Python code to implement the solution from scratch. In the process, you will gain hands-on experience with using popular Python libraries such as Keras to build and train your own neural networks from scratch. By the end of this book, you will have mastered the different neural network architectures and created cutting-edge AI projects in Python that will immediately strengthen your machine learning portfolio. What you will learnLearn various neural network architectures and its advancements in AIMaster deep learning in Python by building and training neural networkMaster neural networks for regression and classificationDiscover convolutional neural networks for image recognitionLearn sentiment analysis on textual data using Long Short-Term MemoryBuild and train a highly accurate facial recognition security systemWho this book is for This book is a perfect match for data scientists, machine learning engineers, and deep learning enthusiasts who wish to create practical neural network projects in Python. Readers should already have some basic knowledge of machine learning and neural networks.

The eight-volume set comprising LNCS volumes 9905-9912 constitutes the refereed proceedings of the 14th European Conference on Computer Vision, ECCV 2016, held in Amsterdam, The Netherlands, in October 2016. The 415 revised papers presented were carefully reviewed and selected from 1480 submissions. The papers cover all aspects of computer vision and pattern recognition such as 3D computer vision; computational photography, sensing and display; face and gesture; low-level vision and image processing; motion and tracking; optimization methods; physics-based vision, photometry and shape-from-X; recognition: detection, categorization, indexing, matching; segmentation, grouping and shape representation; statistical methods and learning; video: events, activities and surveillance; applications. They are organized in topical sections on detection, recognition and retrieval; scene understanding; optimization; image and video processing; learning; action activity and tracking; 3D; and 9 poster sessions.

14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II

Creating Neural Networks with Python

Develop Deep Learning Models for your Natural Language Problems

Sensitivity Analysis for Neural Networks

Transformers for Natural Language Processing

Deep Learning in Natural Language Processing

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.

As book review editor of the IEEE Transactions on Neural Networks, Mohamad Hassoun has had the opportunity to assess the multitude of books on artificial neural networks that have appeared in recent years. Now, in Fundamentals of Artificial Neural Networks, he provides the first systematic account of artificial neural network paradigms by identifying clearly the fundamental concepts and major methodologies underlying most of the current theory and practice employed by neural network researchers. Such a systematic and unified treatment, although sadly lacking in most recent texts on neural networks, makes the subject more accessible to students and practitioners. Here, important results are integrated in order to more fully explain a wide range of existing empirical observations and commonly used heuristics. There are numerous illustrative examples, over 200 end-of-chapter analytical and computer-based problems that will aid in the development of neural network analysis and design skills, and a bibliography of nearly 700 references. Proceeding in a clear and logical fashion, the first two chapters present the basic building blocks and concepts of artificial neural networks and analyze the computational capabilities of the basic network architectures involved. Supervised, reinforcement, and unsupervised learning rules in simple nets are brought together in a common framework in chapter three. The convergence and solution properties of these learning rules are then treated mathematically in chapter four, using the “average learning equation” analysis approach. This organization of material makes it natural to switch into learning multilayer nets using backprop and its variants, described in chapter five. Chapter six covers most of the major neural network paradigms, while associative memories and energy minimizing nets are given detailed coverage in the next chapter. The final chapter takes up Boltzmann machines and Boltzmann learning along with other global search/optimization algorithms such as stochastic gradient search, simulated annealing, and genetic algorithms.

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.

Information in today’s advancing world is rapidly expanding and becoming widely available. This eruption of data has made handling it a daunting and time-consuming task. Natural language processing (NLP) is a method that applies linguistics and algorithms to large amounts of this data to make it more valuable. NLP improves the interaction between humans and computers, yet there remains a lack of research that focuses on the practical implementations of this trending approach. Neural Networks for Natural Language Processing is a collection of innovative research on the methods and applications of linguistic information processing and its computational properties. This publication will support readers with performing sentence classification and language generation using neural networks, apply deep learning models to solve machine translation and conversation problems, and apply deep structured semantic models on information retrieval and natural language applications. While highlighting topics including deep learning, query entity recognition, and information retrieval, this book is ideally designed for research and development professionals, IT specialists, industrialists, technology developers, data analysts, data scientists, academics, researchers, and students seeking current research on the fundamental concepts and techniques of natural language processing.

Introduction to Graph Neural Networks

Neural Representations of Natural Language

Supervised Machine Learning for Text Analysis in R

A Textbook

Neural Network Projects with Python

Methods and Applications

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.

Gain the knowledge of various deep neural network architectures and their application areas to conquer your NLP issues. Key FeaturesGain insights into the basic building blocks of natural language processingLearn how to select the best deep neural network to solve your NLP problemsExplore convolutional and recurrent neural networks and long short-term memory networksBook Description Applying deep learning approaches to various NLP tasks can take your computational algorithms to a completely new level in terms of speed and accuracy. Deep Learning for Natural Language Processing starts off by highlighting the basic building blocks of the natural language processing domain. The book goes on to introduce the problems that you can solve using state-of-the-art neural network models. After this, delving into the various neural network architectures and their specific areas of application will help you to understand how to select the best model to suit your needs. As you advance through this deep learning book, you’ll study convolutional, recurrent, and recursive neural networks, in addition to covering long short-term memory networks (LSTM). Understanding these networks will help you to implement their models using Keras. In the later chapters, you will be able to develop a trigger word detection application using NLP techniques such as attention model and beam search. By the end of this book, you will not only have sound knowledge of natural language processing but also be able to select the best text pre-processing and neural network models to solve a number of NLP issues. What you will learnUnderstand various pre-processing techniques for deep learning problemsBuild a vector representation of text using word2vec and GloVeCreate a named entity recognizer and parts-of-speech tagger with Apache OpenNLPBuild a machine translation model in KerasDevelop a text generation application using LSTMBuild a trigger word detection application using an attention modelWho this book is for If you’re an aspiring data scientist looking for an introduction to deep learning in the NLP domain, this is just the book for you. Strong working knowledge of Python, linear algebra, and machine learning is a must.

Embeddings have undoubtedly been one of the most influential research areas in Natural Language Processing (NLP). Encoding information into a low-dimensional vector representation, which is easily integrable in modern machine learning models, has played a central role in the development of NLP. Embedding techniques initially focused on words, but the attention soon started to shift to other forms: from graph structures, such as knowledge bases, to other types of textual content, such as sentences and documents. This book provides a high-level synthesis of the main embedding techniques in NLP, in the broad sense. The book starts by explaining conventional word vector space models and word embeddings (e.g., Word2Vec and GloVe) and then moves to other types of embeddings, such as word sense, sentence and document, and graph embeddings. The book also provides an overview of recent developments in contextualized representations (e.g., ELMo and BERT) and explains their potential in NLP. Throughout the book, the reader can find both essential information for understanding a certain topic from scratch and a broad overview of the most successful techniques developed in the literature.

Artificial neural networks are used to model systems that receive inputs and produce outputs. The relationships between the inputs and outputs and the representation parameters are critical issues in the design of related engineering systems, and sensitivity analysis concerns methods for analyzing these relationships. Perturbations of neural networks are caused by machine imprecision, and they can be simulated by embedding disturbances in the original inputs or connection weights, allowing us to study the characteristics of a function under small perturbations of its parameters. This is the first book to present a systematic description of sensitivity analysis methods for artificial neural networks. It covers sensitivity analysis of multilayer perceptron neural networks and radial basis function neural networks, two widely used models in the machine learning field. The authors examine the applications of such analysis in tasks such as feature selection, sample reduction, and network optimization. The book will be useful for engineers applying neural network sensitivity analysis to solve practical problems, and for researchers interested in foundational problems in neural networks.

Build innovative deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, RoBERTa, and more

Pattern Recognition and Neural Networks

Neural Networks and Statistical Learning

A Visual, Interactive Guide to Artificial Intelligence

Theory and Advances in Vector Representations of Meaning

Keras 2.x Projects

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You’ll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

This book introduces a variety of neural network methods for solving differential equations arising in science and engineering. The emphasis is placed on a deep understanding of the neural network techniques, which has been presented in a mostly heuristic and intuitive manner. This approach will enable the reader to understand the working, efficiency and shortcomings of each neural network technique for solving differential equations. The objective of this book is to provide the reader with a sound understanding of the foundations of neural networks and a comprehensive introduction to neural network methods for solving differential equations together with recent developments in the techniques and their applications. The book comprises four major sections. Section I consists of a brief overview of differential equations and the relevant physical problems arising in science and engineering. Section II illustrates the history of neural networks starting from their beginnings in the 1940s through to the renewed interest of the 1980s. A general introduction to neural networks and learning technologies is presented in Section III. This section also includes the description of the multilayer perceptron and its learning methods. In Section IV, the different neural network methods for solving differential equations are introduced, including discussion of the most recent developments in the field. Advanced students and researchers in mathematics, computer science and various disciplines in science and engineering will find this book a valuable reference source.

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

Artificial Neural Networks in Biological and Environmental Analysis

Neural Networks for Pattern Recognition

Neural Networks and Deep Learning

Neural Network Methods in Natural Language Processing

The ultimate guide to using Python to explore the true power of neural networks through six projects

7th International Work-Conference on Artificial and Natural Neural Networks, IWANN 2003, Ma ó , Menorca, Spain, June 3-6. Proceedings

Originating from models of biological neural systems, artificial neural networks (ANN) are the cornerstones of artificial intelligence research. Catalyzed by the upsurge in computational power and availability, and made widely accessible with the co-evolution of software, algorithms, and methodologies, artificial neural networks have had a profound impact in the elucidation of complex biological, chemical, and environmental processes. Artificial Neural Networks in Biological and Environmental Analysis provides an in-depth and timely perspective on the fundamental, technological, and applied aspects of computational neural networks. Presenting the basic principles of neural networks together with applications in the field, the book stimulates communication and partnership among scientists in fields as diverse as biology, chemistry, mathematics, medicine, and environmental science. This interdisciplinary discourse is essential not only for the success of independent and collaborative research and teaching programs, but also for the continued interest in the use of neural network tools in scientific inquiry. The book covers: A brief history of computational neural network models in relation to brain function Neural network operations, including neuron connectivity and layer arrangement Basic building blocks of model design, selection, and application from a statistical perspective Neurofuzzy systems, neuro-genetic systems, and neuro-fuzzy-genetic systems Function of neural networks in the study of complex natural processes Scientists deal with very complicated systems, much of the inner workings of which are frequently unknown to researchers. Using only simple, linear mathematical models, information that is needed to truly understand natural systems may be lost. The development of new algorithms to model such processes is needed, and ANNs can play a major role. Balancing basic principles and diverse applications, this text introduces newcomers to the field and reviews recent developments of interest to active neural network practitioners.

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.

Artificial neural networks are nonlinear mapping systems whose structure is loosely based on principles observed in the nervous systems of humans and animals. The basic idea is that massive systems of simple units linked together in appropriate ways can generate many complex and interesting behaviors. This book focuses on the subset of feedforward artificial neural networks called multilayer perceptrons (MLP). These are the mostly widely used neural networks, with applications as diverse as finance (forecasting), manufacturing (process control), and science (speech and image recognition).This book presents an extensive and practical overview of almost every aspect of MLP methodology, progressing from an initial discussion of what MLPs are and how they

might be used to an in-depth examination of technical factors affecting performance. The book can be used as a tool kit by readers interested in applying networks to specific problems, yet it also presents theory and references outlining the last ten years of MLP research.

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

Computer Vision – ECCV 2016

Artificial Neural Nets. Problem Solving Methods

Efficient Processing of Deep Neural Networks

Natural Language Processing with Python

Linguistics for the Age of AI

Deep Learning on Graphs

A human-inspired, linguistically sophisticated model of language understanding for intelligent agent systems. One of the original goals of artificial intelligence research was to endow intelligent agents with human-level natural language capabilities. Recent AI research, however, has focused on applying statistical and machine learning approaches to big data rather than attempting to model what people do and how they do it. In this book, Marjorie McShane and Sergei Nirenburg return to the original goal of recreating human-level intelligence in a machine. They present a human-inspired, linguistically sophisticated model of language understanding for intelligent agent systems that emphasizes meaning—the deep, context-sensitive meaning that a person derives from spoken or written language.

Demonstrate fundamentals of Deep Learning and neural network methodologies using Keras 2.x Key Features Experimental projects showcasing the implementation of high-performance deep learning models with Keras. Use-cases across reinforcement learning, natural language processing, GANs and computer vision. Build strong fundamentals of Keras in the area of deep learning and artificial intelligence. Book Description Keras 2.x Projects explains how to leverage the power of Keras to build and train state-of-the-art deep learning models through a series of practical projects that look at a range of real-world application areas. To begin with, you will quickly set up a deep learning environment by installing the Keras library. Through each of the projects, you will explore and learn the advanced concepts of deep learning and will learn how to compute and run your deep learning models using the advanced offerings of Keras. You will train fully-connected multilayer networks, convolutional neural networks, recurrent neural networks, autoencoders and generative adversarial networks using real-world training datasets. The projects you will undertake are all based on real-world scenarios of all complexity levels, covering topics such as language recognition, stock volatility, energy consumption prediction, faster object classification for self-driving vehicles, and more. By the end of this book, you will be well versed with deep learning and its implementation with Keras. You will have all the knowledge you need to train your own deep learning models to solve different kinds of problems. What you will learn Apply regression methods to your data and understand how the regression algorithm works Understand the basic concepts of classification methods and how to implement them in the Keras environment Import and organize data for neural network classification analysis Learn about the role of rectified linear units in the Keras network architecture Implement a recurrent neural network to classify the sentiment of sentences from movie reviews Set the embedding layer and the tensor sizes of a network Who this book is for If you are a data scientist, machine learning engineer, deep learning practitioner or an AI engineer who wants to build speedy intelligent applications with minimal lines of codes, then this book is the best fit for you. Sound knowledge of machine learning and basic familiarity with Keras library would be useful.

This book provides a broad yet detailed introduction to neural networks and machine learning in a statistical framework. A single, comprehensive resource for study and further research, it explores the major popular neural network models and statistical learning approaches with examples and exercises and allows readers to gain a practical working understanding of the content. This updated new edition presents recently published results and includes six new chapters that correspond to the recent advances in computational learning theory, sparse coding, deep learning, big data and cloud computing. Each chapter features state-of-the-art descriptions and significant research findings. The topics covered include: • multilayer perceptron; • the Hopfield network; • associative memory models; • clustering models and algorithms; • the radial basis function network; • recurrent neural networks; • nonnegative matrix factorization; • independent component analysis; • probabilistic and Bayesian networks; and • fuzzy sets and logic. Focusing on the prominent accomplishments and their practical aspects, this book provides academic and technical staff, as well as graduate students and researchers with a solid foundation and comprehensive reference on the fields of neural networks, pattern recognition, signal processing, and machine learning.

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

An Introduction to Neural Network Methods for Differential Equations

A Systematic Introduction

Graph Representation Learning

Solve your natural language processing problems with smart deep neural networks

Deep Learning Illustrated

Neural Networks

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

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.

Neural networks are a computing paradigm that is finding increasing attention among computer scientists. In this book, theoretical laws and models previously scattered in the literature are brought together into a general theory of artificial neural nets. Always with a view to biology and starting with the simplest nets, it is shown how the properties of models change when more general computing elements and net topologies are introduced. Each chapter contains examples, numerous illustrations, and a bibliography. The book is aimed at readers who seek an overview of the field or who wish to deepen their knowledge. It is suitable as a basis for university courses in neurocomputing.

Graphs are useful data structures in complex real-life applications such as modeling physical systems, learning molecular fingerprints, controlling traffic networks, and recommending friends in social networks. However, these tasks require dealing with non-Euclidean graph data that contains rich relational information between elements and cannot be well handled by traditional deep learning models (e.g., convolutional neural networks (CNNs) or recurrent neural networks (RNNs)). Nodes in graphs usually contain useful feature information that cannot be well addressed in most unsupervised representation learning methods (e.g., network embedding methods). Graph neural networks (GNNs) are proposed to combine the feature information and the graph structure to learn better representations on graphs via feature propagation and aggregation. Due to its convincing performance and high interpretability, GNN has recently become a widely applied graph analysis tool. This book provides a comprehensive introduction to the basic concepts, models, and applications of graph neural networks. It starts with the introduction of the vanilla GNN model. Then several variants of the vanilla model are introduced such as graph convolutional networks, graph recurrent networks, graph attention networks, graph residual networks, and several general frameworks. Variants for different graph types and advanced training methods are also included. As for the applications of GNNs, the book categorizes them into structural, non-structural, and other scenarios, and then it introduces several typical models on solving these tasks. Finally, the closing chapters provide GNN open resources and the outlook of several future directions.

Deep Learning for Natural Language Processing

Supervised Learning in Feedforward Artificial Neural Networks

An Introduction to Neural Networks

Neural Smithing

Neural Network Methods for Natural Language Processing

Analyzing Text with the Natural Language Toolkit

This book offers a highly accessible introduction to natural language processing, the field that supports a variety of language technologies, from predictive text and email filtering to automatic summarization and translation. With it, you'll learn how to write Python programs that work with large collections of unstructured text. You'll access richly annotated datasets using a comprehensive range of linguistic data structures, and you'll understand the main algorithms for analyzing the content and structure of written communication. Packed with examples and exercises, Natural Language Processing with Python will help you: Extract information from unstructured text, either to guess the topic or identify "named entities" Analyze linguistic structure in text, including parsing and semantic analysis Access popular linguistic databases, including WordNet and treebanks Integrate techniques drawn from fields as diverse as linguistics and artificial intelligence This book will help you gain practical skills in natural language processing using the Python programming language and the Natural Language Toolkit (NLTK) open source library. If you're interested in developing web applications, analyzing multilingual news sources, or documenting endangered languages -- or if you're simply curious to have a programmer's perspective on how human language works -- you'll find Natural Language Processing with Python both fascinating and immensely useful.

This book is a collection of chapters describing work carried out as part of a large project at BT Laboratories to study the application of connectionist methods to problems in vision, speech and natural language processing. Also, since the theoretical formulation and the hardware realization of neural networks are significant tasks in themselves, these problems too were addressed. The book, therefore, is divided into five Parts, reporting results in vision, speech, natural language, hardware implementation and network architectures. The three editors of this book have, at one time or another, been involved in planning and running the connectionist project. From the outset, we were concerned to involve the academic community as widely as possible, and consequently, in its first year, over thirty university research groups were funded for small scale studies on the various topics. Co-ordinating such a widely spread project was no small task, and in order to concentrate minds and resources, sets of test problems were devised which were typical of the application areas and were difficult enough to be worthy of study. These are described in the text, and constitute one of the successes of the project.

Publisher's Note: A new edition of this book is out now that includes working with GPT-3 and comparing the results with other models. It includes even more use cases, such as casual language analysis and computer vision tasks, as well as an introduction to OpenAI's Codex. Key FeaturesBuild and implement state-of-the-art language models, such as the original Transformer, BERT, T5, and GPT-2, using concepts that outperform classical deep learning modelsGo through hands-on applications in Python using Google Colaboratory Notebooks with nothing to install on a local machineTest transformer models on advanced use casesBook Description The transformer architecture has proved to be revolutionary in outperforming the classical RNN and CNN models in use today. With an apply-as-you-learn approach, Transformers for Natural Language Processing investigates in vast detail the deep learning for machine translations, speech-to-text, text-to-speech, language modeling, question answering, and many more NLP domains with transformers. The book takes you through NLP with Python and examines various eminent models and datasets within the transformer architecture created by pioneers such as Google, Facebook, Microsoft, OpenAI, and Hugging Face. The book trains you in three stages. The first stage introduces you to transformer architectures, starting with the original transformer, before moving on to RoBERTa, BERT, and DistilBERT models. You will discover training methods for smaller transformers that can outperform GPT-3 in some cases. In the second stage, you will apply transformers for Natural Language Understanding (NLU) and Natural Language Generation (NLG). Finally, the third stage will help you grasp advanced language understanding techniques such as optimizing social network datasets and fake news identification. By the end of this NLP book, you will understand transformers from a cognitive science perspective and be proficient in applying pretrained transformer models by tech giants to various datasets. What you will learnUse the latest pretrained transformer modelsGrasp the workings of the original Transformer, GPT-2, BERT, T5, and other transformer modelsCreate language understanding Python programs using concepts that outperform classical deep learning modelsUse a variety of NLP platforms, including Hugging Face, Trax, and AllenNLPApply Python, TensorFlow, and Keras programs to sentiment analysis, text summarization, speech recognition, machine translations, and moreMeasure the productivity of key transformers to define their scope, potential, and limits in productionWho this book is for Since the book does not teach basic programming, you must be familiar with neural networks, Python, PyTorch, and TensorFlow in order to learn their implementation with Transformers. Readers who can benefit the most from this book include experienced deep learning & NLP practitioners and data analysts & data scientists who want to process the increasing amounts of language-driven data.

This 1996 book explains the statistical framework for pattern recognition and machine learning, now in paperback.

Neural Networks for Vision, Speech and Natural Language

Natural Language Processing

Introduction to Natural Language Processing

A Machine Learning Perspective

Neural Networks for Natural Language Processing

Embeddings in Natural Language Processing

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

'Readers will emerge with a rigorous statistical grounding in the theory of how to construct and train neural networks in pattern recognition' New Scientist

The two-volume set LNCS 2686 and LNCS 2687 constitute the refereed proceedings of the 7th International Work-Conference on Artificial and Natural Neural Networks, IWANN 2003, held in Maó, Menorca, Spain in June 2003. The 197 revised papers presented and selected for inclusion in the book and address the following topics: mathematical and computational methods in neural modelling, neurophysiological data analysis and modelling, structural and functional models of neurons, learning and other plasticity phenomena, dynamics, cognitive processes and artificial intelligence, methodologies for net design, bio-inspired systems and engineering, and applications in a broad variety of fields.

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason and act on this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by these advances.

These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and a practical guide to graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, graph neural networks, and graph neural networks, and applies these methods to walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graphs, and concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning.

9 projects demonstrating faster experimentation of neural network and deep learning applications using Keras

Fundamentals of Artificial Neural Networks

Deep Learning

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn-with visionary content from Grant Beyleveld and beautiful illustrations by Agla é Bassens-presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

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