

Natural Language Processing With Pytorch

Oreilly Media

As recognized, adventure as well as experience virtually lesson, amusement, as skillfully as arrangement can be gotten by just checking out a book **natural language processing with pytorch oreilly media** next it is not directly done, you could agree to even more with reference to this life, something like the world.

We allow you this proper as competently as easy exaggeration to acquire those all. We have the funds for natural language processing with pytorch oreilly media and numerous book collections from fictions to scientific research in any way. among them is this natural language processing with pytorch oreilly media that can be your partner.

Practical Machine Learning for Computer Vision - Valliappa Lakshmanan
2021-07-21

This practical book shows you how to employ machine learning models to extract

information from images. ML engineers and data scientists will learn how to solve a variety of image problems including classification, object detection, autoencoders, image generation, counting,

and captioning with proven ML techniques. This book provides a great introduction to end-to-end deep learning: dataset creation, data preprocessing, model design, model training, evaluation, deployment, and interpretability. Google engineers Valliappa Lakshmanan, Martin Görner, and Ryan Gillard show you how to develop accurate and explainable computer vision ML models and put them into large-scale production using robust ML architecture in a flexible and maintainable way. You'll learn how to design, train, evaluate, and predict with models written in TensorFlow or Keras. You'll learn how to: Design ML architecture for computer vision tasks Select a model (such as ResNet, SqueezeNet, or EfficientNet) appropriate to your task Create an end-to-end ML pipeline to train, evaluate, deploy, and explain your model Preprocess images for data augmentation and to support learnability Incorporate explainability and

responsible AI best practices Deploy image models as web services or on edge devices Monitor and manage ML models

Generative Deep Learning - David Foster
2019-06-28

Generative modeling is one of the hottest topics in AI. It's now possible to teach a machine to excel at human endeavors such as painting, writing, and composing music. With this practical book, machine-learning engineers and data scientists will discover how to re-create some of the most impressive examples of generative deep learning models, such as variational autoencoders, generative adversarial networks (GANs), encoder-decoder models and world models. Author David Foster demonstrates the inner workings of each technique, starting with the basics of deep learning before advancing to some of the most cutting-edge algorithms in the field. Through tips and tricks, you'll understand

how to make your models learn more efficiently and become more creative. Discover how variational autoencoders can change facial expressions in photos Build practical GAN examples from scratch, including CycleGAN for style transfer and MuseGAN for music generation Create recurrent generative models for text generation and learn how to improve the models using attention Understand how generative models can help agents to accomplish tasks within a reinforcement learning setting Explore the architecture of the Transformer (BERT, GPT-2) and image generation models such as ProGAN and StyleGAN

Mastering Machine Learning with

scikit-learn - Gavin Hackeling 2017-07-24

Use scikit-learn to apply machine learning to real-world problems About This Book Master popular machine learning models including k-nearest neighbors, random forests, logistic

regression, k-means, naive Bayes, and artificial neural networks Learn how to build and evaluate performance of efficient models using scikit-learn Practical guide to master your basics and learn from real life applications of machine learning Who This Book Is For This book is intended for software engineers who want to understand how common machine learning algorithms work and develop an intuition for how to use them, and for data scientists who want to learn about the scikit-learn API. Familiarity with machine learning fundamentals and Python are helpful, but not required. What You Will Learn Review fundamental concepts such as bias and variance Extract features from categorical variables, text, and images Predict the values of continuous variables using linear regression and K Nearest Neighbors Classify documents and images using logistic regression and support vector machines Create ensembles of

estimators using bagging and boosting techniques Discover hidden structures in data using K-Means clustering Evaluate the performance of machine learning systems in common tasks In Detail Machine learning is the buzzword bringing computer science and statistics together to build smart and efficient models. Using powerful algorithms and techniques offered by machine learning you can automate any analytical model. This book examines a variety of machine learning models including popular machine learning algorithms such as k-nearest neighbors, logistic regression, naive Bayes, k-means, decision trees, and artificial neural networks. It discusses data preprocessing, hyperparameter optimization, and ensemble methods. You will build systems that classify documents, recognize images, detect ads, and more. You will learn to use scikit-learn's API to extract features from categorical variables, text and images; evaluate model

performance, and develop an intuition for how to improve your model's performance. By the end of this book, you will master all required concepts of scikit-learn to build efficient models at work to carry out advanced tasks with the practical approach. Style and approach This book is motivated by the belief that you do not understand something until you can describe it simply. Work through toy problems to develop your understanding of the learning algorithms and models, then apply your learnings to real-life problems.

Deep Learning with PyTorch - Luca Pietro Giovanni Antiga 2020-07-01

“We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch’s creator and key contributors

Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This

practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in

machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 -

LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production *Natural Language Processing with Python* - Steven Bird 2009-06-12 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. *Natural Language Processing with Python Quick Start Guide* - Nirant Kasliwal
2018-11-30

Build and deploy intelligent applications for natural language processing with Python by using industry standard tools and recently popular methods in deep learning Key FeaturesA no-math, code-driven programmer's guide to text processing and NLPGet state of the art results with modern tooling across linguistics, text vectors and machine learningFundamentals of NLP methods from spaCy, gensim, scikit-learn and PyTorchBook Description NLP in Python is among the most sought after skills among data scientists. With code and relevant case studies, this book will show how you can use industry-grade tools to implement NLP

programs capable of learning from relevant data. We will explore many modern methods ranging from spaCy to word vectors that have reinvented NLP. The book takes you from the basics of NLP to building text processing applications. We start with an introduction to the basic vocabulary along with a workflow for building NLP applications. We use industry-grade NLP tools for cleaning and pre-processing text, automatic question and answer generation using linguistics, text embedding, text classifier, and building a chatbot. With each project, you will learn a new concept of NLP. You will learn about entity recognition, part of speech tagging and dependency parsing for Q and A. We use text embedding for both clustering documents and making chatbots, and then build classifiers using scikit-learn. We conclude by deploying these models as REST APIs with Flask. By the end, you will be confident building NLP

applications, and know exactly what to look for when approaching new challenges. What you will learn

Understand classical linguistics in using English grammar for automatically generating questions and answers from a free text corpus

Work with text embedding models for dense number representations of words, subwords and characters in the English language for exploring document clustering

Deep Learning in NLP using PyTorch with a code-driven introduction to PyTorch

Using an NLP project management Framework for estimating timelines and organizing your project into stages

Hack and build a simple chatbot application in 30 minutes

Deploy an NLP or machine learning application using Flask as RESTFUL APIs

Who this book is for

Programmers who wish to build systems that can interpret language. Exposure to Python programming is required. Familiarity with NLP or machine learning vocabulary will be helpful, but not

mandatory.

Applied Natural Language Processing

in the Enterprise - Ankur A. Patel

2021-05-12

NLP has exploded in popularity over the last few years. But while Google, Facebook, OpenAI, and others continue to release larger language models, many teams still struggle with building NLP applications that live up to the hype. This hands-on guide helps you get up to speed on the latest and most promising trends in NLP. With a basic understanding of machine learning and some Python experience, you'll learn how to build, train, and deploy models for real-world applications in your organization.

Authors Ankur Patel and Ajay Uppili

Arasanipalai guide you through the process using code and examples that highlight the best practices in modern NLP. Use state-of-the-art NLP models such as BERT and GPT-3 to solve NLP tasks such as named entity

recognition, text classification, semantic search, and reading comprehension Train NLP models with performance comparable or superior to that of out-of-the-box systems Learn about Transformer architecture and modern tricks like transfer learning that have taken the NLP world by storm Become familiar with the tools of the trade, including spaCy, Hugging Face, and fast.ai Build core parts of the NLP pipeline--including tokenizers, embeddings, and language models--from scratch using Python and PyTorch Take your models out of Jupyter notebooks and learn how to deploy, monitor, and maintain them in production

Practical Weak Supervision

- Wee Hyong Tok 2021-09-30

Most data scientists and engineers today rely on quality labeled data to train machine learning models. But building a training set manually is time-consuming and expensive, leaving many companies with unfinished ML

projects. There's a more practical approach. In this book, Wee Hyong Tok, Amit Bahree, and Senja Filipi show you how to create products using weakly supervised learning models. You'll learn how to build natural language processing and computer vision projects using weakly labeled datasets from Snorkel, a spin-off from the Stanford AI Lab. Because so many companies have pursued ML projects that never go beyond their labs, this book also provides a guide on how to ship the deep learning models you build. Get up to speed on the field of weak supervision, including ways to use it as part of the data science process Use Snorkel AI for weak supervision and data programming Get code examples for using Snorkel to label text and image datasets Use a weakly labeled dataset for text and image classification Learn practical considerations for using Snorkel with large datasets and using Spark clusters to scale labeling

[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 Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book 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 learn

- Cover advanced and state-of-the-art neural network architectures
- Understand the theory and math behind neural networks
- Train DNNs and apply them to modern deep learning problems
- Use CNNs for object detection and image segmentation
- Implement generative adversarial networks (GANs) and variational autoencoders to generate new images
- Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models
- Understand DL techniques, such as meta-learning and graph neural networks

Who 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.

Deep Learning with Python - Francois Chollet 2017-11-30

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples.

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image

recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from

first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the

International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Natural Language Processing with Spark NLP - Alex Thomas 2020-06-25

If you want to build an enterprise-quality application that uses natural language text but aren't sure where to begin or what tools to use, this practical guide will help get you started. Alex Thomas, principal data

scientist at Wisecube, shows software engineers and data scientists how to build scalable natural language processing (NLP) applications using deep learning and the Apache Spark NLP library. Through concrete examples, practical and theoretical explanations, and hands-on exercises for using NLP on the Spark processing framework, this book teaches you everything from basic linguistics and writing systems to sentiment analysis and search engines. You'll also explore special concerns for developing text-based applications, such as performance. In four sections, you'll learn NLP basics and building blocks before diving into application and system building: Basics: Understand the fundamentals of natural language processing, NLP on Apache Spark, and deep learning Building blocks: Learn techniques for building NLP applications—including tokenization, sentence segmentation, and named-entity

recognition—and discover how and why they work Applications: Explore the design, development, and experimentation process for building your own NLP applications Building NLP systems: Consider options for productionizing and deploying NLP models, including which human languages to support

Natural Language Processing in Action

- Hannes Hapke 2019-03-16

Summary Natural Language Processing in Action is your guide to creating machines that understand human language using the power of Python with its ecosystem of packages dedicated to NLP and AI. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Recent advances in deep learning empower applications to understand text and speech with extreme accuracy. The result? Chatbots that can imitate real people, meaningful

resume-to-job matches, superb predictive search, and automatically generated document summaries—all at a low cost. New techniques, along with accessible tools like Keras and TensorFlow, make professional-quality NLP easier than ever before. About the Book Natural Language Processing in Action is your guide to building machines that can read and interpret human language. In it, you'll use readily available Python packages to capture the meaning in text and react accordingly. The book expands traditional NLP approaches to include neural networks, modern deep learning algorithms, and generative techniques as you tackle real-world problems like extracting dates and names, composing text, and answering free-form questions. What's inside Some sentences in this book were written by NLP! Can you guess which ones? Working with Keras, TensorFlow, gensim, and scikit-learn Rule-

based and data-based NLP Scalable pipelines About the Reader This book requires a basic understanding of deep learning and intermediate Python skills. About the Author Hobson Lane, Cole Howard, and Hannes Max Hapke are experienced NLP engineers who use these techniques in production. Table of Contents PART 1 - WORDY MACHINES Packets of thought (NLP overview) Build your vocabulary (word tokenization) Math with words (TF-IDF vectors) Finding meaning in word counts (semantic analysis) PART 2 - DEEPER LEARNING (NEURAL NETWORKS) Baby steps with neural networks (perceptrons and backpropagation) Reasoning with word vectors (Word2vec) Getting words in order with convolutional neural networks (CNNs) Loopy (recurrent) neural networks (RNNs) Improving retention with long short-term memory networks Sequence-to-sequence models and attention

PART 3 - GETTING REAL (REAL-WORLD NLP CHALLENGES) Information extraction (named entity extraction and question answering) Getting chatty (dialog engines) Scaling up (optimization, parallelization, and batch processing)

Deep Learning for Coders with fastai and PyTorch - Jeremy Howard 2020-06-29

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also

dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

AI and Machine Learning for Coders - Laurence Moroney 2020-10-01

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on,

code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

TinyML - Pete Warden 2019-12-16
Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the

essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size
[PyTorch Pocket Reference](#) - Joe Papa
2021-05-11

This concise, easy-to-use reference puts one of the most popular frameworks for deep learning research and development at your fingertips. Author Joe Papa provides instant access to syntax, design patterns, and code examples to accelerate your development and reduce the time you spend searching for answers. Research scientists, machine learning engineers, and software developers will find clear, structured PyTorch code that covers every step of neural network development—from loading data to

customizing training loops to model optimization and GPU/TPU acceleration. Quickly learn how to deploy your code to production using AWS, Google Cloud, or Azure and deploy your ML models to mobile and edge devices. Learn basic PyTorch syntax and design patterns Create custom models and data transforms Train and deploy models using a GPU and TPU Train and test a deep learning classifier Accelerate training using optimization and distributed training Access useful PyTorch libraries and the PyTorch ecosystem

Deep Learning - Ian Goodfellow
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.

Speech & Language Processing - Dan Jurafsky 2000-09

[Representation Learning for Natural Language Processing](#) - Zhiyuan Liu 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.

Transformers for Natural Language Processing - Denis Rothman 2021-01-29

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 learn
Use the latest pretrained transformer models
Grasp the workings of the original Transformer, GPT-2, BERT, T5, and other transformer models
Create language understanding Python programs using concepts that outperform classical deep learning models
Use a variety of NLP platforms, including Hugging Face, Trax, and AllenNLP
Apply Python, TensorFlow, and Keras programs to sentiment analysis, text summarization, speech recognition, machine translations, and more
Measure the productivity of key transformers to define their scope, potential, and limits in production
Who 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.

Learning Deep Learning - Magnus Ekman
2021-08

NVIDIA's Full-Color Guide to Deep Learning: All Students Need to Get Started and Get Results
Learning Deep Learning is a complete guide to DL. Illuminating both the core concepts and the hands-on programming techniques needed to succeed, this book suits seasoned developers, data scientists, analysts, but also those with no prior machine learning or statistic experience. After introducing the essential building blocks of deep neural networks, such as artificial neurons and fully connected, convolutional, and recurrent layers, Magnus Ekman shows how to use them to build advanced architectures, including the Transformer. He describes how these concepts are used to build

modern networks for computer vision and natural language processing (NLP), including Mask R-CNN, GPT, and BERT. And he explains how a natural language translator and a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples using TensorFlow with Keras. Corresponding PyTorch examples are provided online, and the book thereby covers the two dominating Python libraries for DL used in industry and academia. He concludes with an introduction to neural architecture search (NAS), exploring important ethical issues and providing resources for further learning. Explore and master core concepts: perceptrons, gradient-based learning, sigmoid neurons, and back propagation. See how DL frameworks make it easier to develop more complicated and useful neural networks. Discover how convolutional

neural networks (CNNs) revolutionize image classification and analysis. Apply recurrent neural networks (RNNs) and long short-term memory (LSTM) to text and other variable-length sequences. Master NLP with sequence-to-sequence networks and the Transformer architecture. Build applications for natural language translation and image captioning.

Natural Language Processing Mit PyTorch - Delip Rao 2019

Sprachanwendungen wie Amazon Alexa und Google Translate sind heute allgegenwärtig. Grundlage dafür ist das Natural Language Processing (NLP), das zahllose Möglichkeiten für die Entwicklung intelligenter, Deep-Learning-basierter Anwendungen eröffnet. In diesem Buch lernen Sie die neuesten Techniken zur Verarbeitung von Sprache kennen und nutzen dabei das neue, flexible Deep-Learning-Framework PyTorch. Die Autoren vermitteln Ihnen einen Überblick

über NLP-Methoden und Grundkonzepte neuronaler Netze und demonstrieren Ihnen dann, wie Sie Sprachanwendungen mit PyTorch entwickeln. Sie erfahren z.B., wie Sie Einbettungen verwenden, um Wörter, Sätze und Dokumente darzustellen, wie sich Sequenzdaten mit RNNs modellieren und Sequenzvoraussagen und Sequenz-zu-Sequenz-Modelle generieren lassen, und Sie lernen Entwurfsmuster für den Aufbau von produktionsreifen NLP-Systemen kennen.

Frontiers of Quality Electronic Design (QED)

- Ali Iranmanesh 2023-01-11

Quality Electronic Design (QED)'s landscape spans a vast region where territories of many participating disciplines and technologies overlap. This book explores the latest trends in several key topics related to quality electronic design, with emphasis on Hardware Security, Cybersecurity, Machine Learning, and application of Artificial Intelligence (AI). The book includes topics in

nonvolatile memories (NVM), Internet of Things (IoT), FPGA, and Neural Networks.

Applied Text Analysis with Python -

Benjamin Bengfort 2018-06-11

From news and speeches to informal chatter on social media, natural language is one of the richest and most underutilized sources of data. Not only does it come in a constant stream, always changing and adapting in context; it also contains information that is not conveyed by traditional data sources. The key to unlocking natural language is through the creative application of text analytics. This practical book presents a data scientist's approach to building language-aware products with applied machine learning. You'll learn robust, repeatable, and scalable techniques for text analysis with Python, including contextual and linguistic feature engineering, vectorization, classification, topic modeling, entity resolution, graph analysis, and visual

steering. By the end of the book, you'll be equipped with practical methods to solve any number of complex real-world problems. Preprocess and vectorize text into high-dimensional feature representations Perform document classification and topic modeling Steer the model selection process with visual diagnostics Extract key phrases, named entities, and graph structures to reason about data in text Build a dialog framework to enable chatbots and language-driven interaction Use Spark to scale processing power and neural networks to scale model complexity

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow -

Aurélien Géron 2019-09-05

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to

implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net

architectures, including convolutional nets, recurrent nets, and deep reinforcement learning. Learn techniques for training and scaling deep neural nets.

Practical Machine Learning with Python -

Dipanjan Sarkar 2017-12-20

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples

and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing,

computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including

CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students
Programming PyTorch for Deep Learning - Ian Pointer 2019-09-20

Deep learning is changing everything. This machine-learning method has already surpassed traditional computer vision techniques, and the same is happening with NLP. If you're looking to bring deep learning into your domain, this practical book will bring you up to speed on key concepts using Facebook's PyTorch framework. Once author Ian Pointer helps you set up PyTorch on a cloud-based environment, you'll learn how use the framework to create neural architectures for performing operations on images, sound, text, and other types of data. By the end of the book, you'll be able to create neural networks and train them on multiple types of data. Learn how to deploy

deep learning models to production Explore PyTorch use cases in companies other than Facebook Learn how to apply transfer learning to images Apply cutting-edge NLP techniques using a model trained on Wikipedia

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.

Deep Learning Illustrated - Jon Krohn
2019-08-05

"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 Aglaé 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.

The SAGE Handbook of Social Media Research Methods - Anabel Quan-Haase
2022-09-02

The SAGE Handbook of Social Media Research Methods spans the entire research process, from data collection to analysis and

interpretation. This second edition has been comprehensively updated and expanded, from 39 to 49 chapters. In addition to a new section of chapters focussing on ethics, privacy and the politics of social media data, the new edition provides broader coverage of topics such as: Data sources Scraping and spidering data Locative data, video data and linked data Platform-specific analysis Analytical tools Critical social media analysis Written by leading scholars from across the globe, the chapters provide a mix of theoretical and applied assessments of topics, and include a range of new case studies and data sets that exemplify the methodological approaches. This Handbook is an essential resource for any researcher or postgraduate student embarking on a social media research project. PART 1: Conceptualising and Designing Social Media Research PART 2: Collecting Data PART 3: Qualitative Approaches to Social Media Data

PART 4: Quantitative Approaches to Social Media Data PART 5: Diverse Approaches to Social Media Data PART 6: Research & Analytical Tools PART 7: Social Media Platforms PART 8: Privacy, Ethics and Inequalities

Machine Learning Pocket Reference - Matt Harrison 2019-08-27

With detailed notes, tables, and examples, this handy reference will help you navigate the basics of structured machine learning. Author Matt Harrison delivers a valuable guide that you can use for additional support during training and as a convenient resource when you dive into your next machine learning project. Ideal for programmers, data scientists, and AI engineers, this book includes an overview of the machine learning process and walks you through classification with structured data. You'll also learn methods for clustering, predicting a continuous value (regression),

and reducing dimensionality, among other topics. This pocket reference includes sections that cover: Classification, using the Titanic dataset Cleaning data and dealing with missing data Exploratory data analysis Common preprocessing steps using sample data Selecting features useful to the model Model selection Metrics and classification evaluation Regression examples using k-nearest neighbor, decision trees, boosting, and more Metrics for regression evaluation Clustering Dimensionality reduction Scikit-learn pipelines

Natural Language Processing with PyTorch - Delip Rao 2019-01-22

Natural Language Processing (NLP) provides boundless opportunities for solving problems in artificial intelligence, making products such as Amazon Alexa and Google Translate possible. If you're a developer or data scientist new to NLP and deep learning, this practical guide shows you how to apply

these methods using PyTorch, a Python-based deep learning library. Authors Delip Rao and Brian McMahon provide you with a solid grounding in NLP and deep learning algorithms and demonstrate how to use PyTorch to build applications involving rich representations of text specific to the problems you face. Each chapter includes several code examples and illustrations. Explore computational graphs and the supervised learning paradigm Master the basics of the PyTorch optimized tensor manipulation library Get an overview of traditional NLP concepts and methods Learn the basic ideas involved in building neural networks Use embeddings to represent words, sentences, documents, and other features Explore sequence prediction and generate sequence-to-sequence models Learn design patterns for building production NLP systems [Advances in Machine Learning/Deep](#)

Learning-based Technologies - George A. Tsihrintzis 2021-08-05

As the 4th Industrial Revolution is restructuring human societal organization into, so-called, "Society 5.0", the field of Machine Learning (and its sub-field of Deep Learning) and related technologies is growing continuously and rapidly, developing in both itself and towards applications in many other disciplines. Researchers worldwide aim at incorporating cognitive abilities into machines, such as learning and problem solving. When machines and software systems have been enhanced with Machine Learning/Deep Learning components, they become better and more efficient at performing specific tasks. Consequently, Machine Learning/Deep Learning stands out as a research discipline due to its worldwide pace of growth in both theoretical advances and areas of application, while achieving

very high rates of success and promising major impact in science, technology and society. The book at hand aims at exposing its readers to some of the most significant Advances in Machine Learning/Deep Learning-based Technologies. The book consists of an editorial note and an additional ten (10) chapters, all invited from authors who work on the corresponding chapter theme and are recognized for their significant research contributions. In more detail, the chapters in the book are organized into five parts, namely (i) Machine Learning/Deep Learning in Socializing and Entertainment, (ii) Machine Learning/Deep Learning in Education, (iii) Machine Learning/Deep Learning in Security, (iv) Machine Learning/Deep Learning in Time Series Forecasting, and (v) Machine Learning in Video Coding and Information Extraction. This research book is directed towards professors, researchers, scientists,

engineers and students in Machine Learning/Deep Learning-related disciplines. It is also directed towards readers who come from other disciplines and are interested in becoming versed in some of the most recent Machine Learning/Deep Learning-based technologies. An extensive list of bibliographic references at the end of each chapter guides the readers to probe further into the application areas of interest to them.

Applied Natural Language Processing in the Enterprise - Ankur Patel 2021

NLP is one of the hottest topics in AI today. Having lagged for years behind other deep learning fields such as computer vision, NLP only recently gained mainstream popularity. Google, Facebook, and OpenAI have open-sourced large pretrained language models, but many organizations today still struggle with building and adopting NLP applications. This hands-on guide helps you learn the

process quickly. If you have a basic to intermediate understanding of machine learning and programming experience with Python, you'll learn how to build and deploy real-world NLP applications in your organization. Authors Ankur Patel and Ajay Uppili Arasanipalai walk you through the process without bogging you down in theory. Understand how state-of-the-art NLP models work Learn the tools of the trade, including frameworks popular today Perform NLP tasks such as text classification, semantic search, and reading comprehension Solve problems using new models like transformers and techniques such as transfer learning Build NLP models from scratch with performance comparable or superior to out-of-the-box systems Deploy your models to production and maintain their performance Implement a suite of NLP algorithms using Python and PyTorch.

Applied Natural Language Processing

in the Enterprise - Ankur A. Patel

2021-05-12

NLP has exploded in popularity over the last few years. But while Google, Facebook, OpenAI, and others continue to release larger language models, many teams still struggle with building NLP applications that live up to the hype. This hands-on guide helps you get up to speed on the latest and most promising trends in NLP. With a basic understanding of machine learning and some Python experience, you'll learn how to build, train, and deploy models for real-world applications in your organization. Authors Ankur Patel and Ajay Uppili Arasanipalai guide you through the process using code and examples that highlight the best practices in modern NLP. Use state-of-the-art NLP models such as BERT and GPT-3 to solve NLP tasks such as named entity recognition, text classification, semantic search, and reading comprehension Train

NLP models with performance comparable or superior to that of out-of-the-box systems Learn about Transformer architecture and modern tricks like transfer learning that have taken the NLP world by storm Become familiar with the tools of the trade, including spaCy, Hugging Face, and fast.ai Build core parts of the NLP pipeline--including tokenizers, embeddings, and language models--from scratch using Python and PyTorch Take your models out of Jupyter notebooks and learn how to deploy, monitor, and maintain them in production *Deep Learning for Coders with fastai and PyTorch* - Jeremy Howard 2020-06-29 Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and

minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Machine Learning Design Patterns -

Valliappa Lakshmanan 2020-10-15

The design patterns in this book capture best practices and solutions to recurring problems in machine learning. The authors, three Google engineers, catalog proven methods to help data scientists tackle common problems throughout the ML process. These design patterns codify the experience of hundreds of experts into straightforward, approachable advice. In this book, you will find detailed explanations of 30 patterns for data and problem representation, operationalization, repeatability, reproducibility, flexibility, explainability, and fairness. Each pattern includes a description of the problem, a variety of potential solutions, and recommendations for choosing the best technique for your situation. You'll learn how to: Identify and mitigate common challenges when training, evaluating, and deploying ML

models Represent data for different ML model types, including embeddings, feature crosses, and more Choose the right model type for specific problems Build a robust training loop that uses checkpoints, distribution strategy, and hyperparameter tuning Deploy scalable ML systems that you can retrain and update to reflect new data Interpret model predictions for stakeholders and ensure models are treating users fairly

Natural Language Processing with Transformers - Lewis Tunstall 2022-01-26

Since their introduction in 2017, transformers have quickly become the dominant architecture for achieving state-of-the-art results on a variety of natural language processing tasks. If you're a data scientist or coder, this practical book shows you how to train and scale these large models using Hugging Face Transformers, a Python-based deep learning library. Transformers have been used to write

realistic news stories, improve Google Search queries, and even create chatbots that tell corny jokes. In this guide, authors Lewis Tunstall, Leandro von Werra, and Thomas Wolf, among the creators of Hugging Face Transformers, use a hands-on approach to teach you how transformers work and how to integrate them in your applications. You'll quickly learn a variety of tasks they can help you solve. Build, debug, and optimize transformer models for core NLP tasks, such as text classification, named entity recognition, and question answering Learn how transformers can be used for cross-lingual transfer learning Apply transformers in real-world scenarios where labeled data is scarce Make transformer models efficient for deployment using techniques such as distillation, pruning, and quantization Train transformers from scratch and learn how to scale to multiple GPUs and distributed environments

Practical Natural Language Processing -

Sowmya Vajjala 2020-06-17

Many books and courses tackle natural language processing (NLP) problems with toy use cases and well-defined datasets. But if you want to build, iterate, and scale NLP systems in a business setting and tailor them for particular industry verticals, this is your guide. Software engineers and data scientists will learn how to navigate the maze of options available at each step of the journey. Through the course of the book, authors Sowmya Vajjala, Bodhisattwa Majumder, Anuj Gupta, and Harshit Surana will guide you through the process of building real-world NLP solutions embedded in larger product setups. You'll learn how to adapt your solutions for different industry verticals such as healthcare, social media, and retail. With this book, you'll: Understand the wide spectrum of problem statements, tasks, and solution approaches within NLP

Implement and evaluate different NLP applications using machine learning and deep learning methods Fine-tune your NLP solution based on your business problem and industry vertical Evaluate various algorithms and approaches for NLP product tasks, datasets, and stages Produce software solutions following best practices around release, deployment, and DevOps for NLP systems Understand best practices, opportunities, and the roadmap for NLP from a business and product leader's perspective

Data Science on AWS -

Chris Fregly 2021-04-07

With this practical book, AI and machine learning practitioners will learn how to successfully build and deploy data science projects on Amazon Web Services. The Amazon AI and machine learning stack unifies data science, data engineering, and application development to help level up your skills. This guide shows you how to

build and run pipelines in the cloud, then integrate the results into applications in minutes instead of days. Throughout the book, authors Chris Fregly and Antje Barth demonstrate how to reduce cost and improve performance. Apply the Amazon AI and ML stack to real-world use cases for natural language processing, computer vision, fraud detection, conversational devices, and more Use automated machine learning to implement a specific subset of use cases with SageMaker Autopilot Dive deep into the complete model development

lifecycle for a BERT-based NLP use case including data ingestion, analysis, model training, and deployment Tie everything together into a repeatable machine learning operations pipeline Explore real-time ML, anomaly detection, and streaming analytics on data streams with Amazon Kinesis and Managed Streaming for Apache Kafka Learn security best practices for data science projects and workflows including identity and access management, authentication, authorization, and more