

Natural Language Processing With Pytorch O'Reilly Media

Get well-versed with traditional as well as modern natural language processing concepts and techniques

Key Features

- Perform various NLP tasks to build linguistic applications using Python libraries
- Understand, analyze, and generate text to provide accurate results
- Interpret human language using various NLP concepts, methodologies, and tools

Book Description

Natural Language Processing (NLP) is the subfield in computational linguistics that enables computers to understand, process, and analyze text. This book caters to the unmet demand for hands-on training of NLP concepts and provides exposure to real-world applications along with a solid theoretical grounding. This book starts by introducing you to the field of NLP and its applications, along with the modern Python libraries that you'll use to build your NLP-powered apps. With the help of practical examples, you'll learn how to build reasonably sophisticated NLP applications, and cover various methodologies and challenges in deploying NLP applications in the real world. You'll cover key NLP tasks such as text classification, semantic embedding, sentiment analysis, machine translation, and developing a chatbot using machine learning and deep learning techniques. The book will also help you discover how machine learning techniques play a vital role in making your linguistic apps smart. Every chapter is accompanied by examples of real-world applications to help you build impressive NLP applications of your own. By the end of this NLP book, you'll be able to work with language data, use machine learning to identify patterns in text, and get acquainted with the advancements in NLP. What you will learn

- Understand how NLP powers modern applications
- Explore key NLP techniques to build your natural language vocabulary
- Transform text data into mathematical data structures and learn how to improve text mining models
- Discover how various neural network architectures work with natural language data
- Get the hang of building sophisticated text processing models using machine learning and deep learning
- Check out state-of-the-art architectures that have revolutionized research in the NLP domain

Who this book is for

This NLP Python book is for anyone looking to learn NLP's theoretical and practical aspects alike. It starts with the basics and gradually covers advanced concepts to make it easy to follow for readers with varying levels of NLP proficiency. This comprehensive guide will help you develop a thorough understanding of the NLP methodologies for building linguistic applications; however, working knowledge of Python programming language and high school level mathematics is expected.

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.

Get to grips with deep learning techniques for building image processing applications using PyTorch with the help of code notebooks and test questions
Key Features
Implement solutions to 50 real-world computer vision applications using PyTorch
Understand the theory and working mechanisms of neural network architectures and their implementation
Discover best practices using a custom library created especially for this book
Book Description
Deep learning is the driving force behind many recent advances in various computer vision (CV) applications. This book takes a hands-on approach to help you to solve over 50 CV problems using PyTorch 1.x on real-world datasets. You'll start by building a neural network (NN) from scratch using NumPy and PyTorch and discover best practices for tweaking its hyperparameters. You'll then perform image classification using convolutional neural networks and transfer learning and understand how they work. As you progress, you'll implement multiple use cases of 2D and 3D multi-object detection, segmentation, human-pose-estimation by learning about the R-CNN family, SSD, YOLO, U-Net architectures, and the Detectron2 platform. The book will also guide you in performing facial expression swapping, generating new faces, and manipulating facial expressions as you explore autoencoders and modern generative adversarial networks. You'll learn how to combine CV with NLP techniques, such as LSTM and transformer, and RL techniques, such as Deep Q-learning, to implement OCR, image captioning, object detection, and a self-driving car agent. Finally, you'll move your NN model to production on the AWS Cloud. By the end of this book, you'll be able to leverage modern NN architectures to solve over 50 real-world CV problems confidently. What you will learn
Train a NN from scratch with NumPy and PyTorch
Implement 2D and 3D multi-object detection and segmentation
Generate digits and DeepFakes with autoencoders and advanced GANs
Manipulate images using CycleGAN, Pix2PixGAN, StyleGAN2, and SRGAN
Combine CV with NLP to perform OCR, image captioning, and object detection
Combine CV with reinforcement learning to build agents that play pong and self-drive a car
Deploy a deep learning model on the AWS server using FastAPI and Docker
Implement over 35 NN architectures and common OpenCV utilities
Who this book is for
This book is for beginners to PyTorch and intermediate-level machine learning practitioners who are looking to get well-versed with computer vision techniques using deep learning and PyTorch. If you are just getting started with neural networks, you'll find the use cases accompanied by notebooks in GitHub present in this book useful. Basic knowledge of the Python programming language and machine learning is all you need to get started with this book.

Turning text into valuable information is essential for businesses looking to gain a competitive advantage. With recent improvements in natural language processing (NLP), users now have many options for solving complex challenges. But it's not always clear

which NLP tools or libraries would work for a business's needs, or which techniques you should use and in what order. This practical book provides data scientists and developers with blueprints for best practice solutions to common tasks in text analytics and natural language processing. Authors Jens Albrecht, Sidharth Ramachandran, and Christian Winkler provide real-world case studies and detailed code examples in Python to help you get started quickly. Extract data from APIs and web pages Prepare textual data for statistical analysis and machine learning Use machine learning for classification, topic modeling, and summarization Explain AI models and classification results Explore and visualize semantic similarities with word embeddings Identify customer sentiment in product reviews Create a knowledge graph based on named entities and their relations

Unlocking Text Data with Machine Learning and Deep Learning using Python

Using Natural Language Processing and Machine Learning
Explore tools and techniques to analyze and process text with a view to building real-world NLP applications

Natural Language Processing with Spark NLP

Build Intelligent Language Applications Using Deep Learning

A Visual, Interactive Guide to Artificial Intelligence

Learning to Understand Text at Scale

Learn how to redesign NLP applications from scratch. KEY FEATURES

- Get familiar with the basics of any Machine Learning or Deep Learning application.
- Understand how does preprocessing work in NLP pipeline.
- Use simple PyTorch snippets to create basic building blocks of the network commonly used in NLP.
- Learn how to build a complex NLP application.

- Get familiar with the advanced embedding technique, Generative network, and Audio signal processing techniques.

DESCRIPTION Natural language processing (NLP) is one of the areas where many Machine Learning and Deep Learning techniques are applied. This book covers wide areas, including the fundamentals of Machine Learning, Understanding and optimizing Hyperparameters, Convolution Neural Networks (CNN), and Recurrent Neural Networks (RNN). This book not only covers the classical concept of text processing but also shares the recent advancements. This book will empower users in designing networks with the least computational and time complexity. This book not only covers basics of Natural Language Processing but also helps in deciphering the logic behind advanced concepts/architecture such as Batch Normalization, Position Embedding, DenseNet, Attention Mechanism, Highway Networks, Transformer models and Siamese Networks. This book also covers recent advancements such as ELMo-BiLM, SkipThought, and Bert. This book also covers practical implementation with step by step explanation of deep learning techniques in Topic Modelling, Text Generation, Named Entity Recognition, Text Summarization, and Language Translation. In addition to this, very advanced and open to research topics such as Generative Adversarial Network and Speech Processing are also covered. *WHAT YOU WILL LEARN*

- Learn how to leveraging GPU for Deep Learning
- Learn how to use complex embedding models such as BERT
- Get familiar with the common NLP applications.
- Learn how to use GANs in NLP
- Learn how to process Speech data and implementing it in Speech applications

WHO THIS BOOK IS FOR This book is a must-read to everyone who wishes to start the career with Machine learning and Deep Learning. This book is also for those who want to use GPU for developing Deep Learning applications. *TABLE OF CONTENTS*

1. Understanding the basics of learning Process
2. Text Processing Techniques
3. Representing Language Mathematically
- 4.

Using RNN for NLP 5. Applying CNN In NLP Tasks 6. Accelerating NLP with Advanced Embeddings 7. Applying Deep Learning to NLP tasks 8. Application of Complex Architectures in NLP 9. Understanding Generative Networks 10. Techniques of Speech Processing 11. The Road Ahead

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 Features
Build 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 models
Go through hands-on applications in Python using Google Colaboratory Notebooks with nothing to install on a local machine
Test transformer models on advanced use cases
Book 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.

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

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

Transfer Learning for Natural Language Processing

Programming PyTorch for Deep Learning

Blueprints for Text Analytics Using Python

Build and train state-of-the-art natural language processing models using BERT

Build effective real-world NLP applications using NER, RNNs, seq2seq models, Transformers, and more

Modern Computer Vision with PyTorch

Build state-of-the-art models from scratch with advanced natural language processing techniques

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. 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

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

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 Stark, 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

Practical Natural Language Processing

Going from a Python developer to an effective Natural Language Processing Engineer

Learn how to build NLP applications with Deep Learning (English Edition)

Understanding, analyzing, and generating text with Python

Build Innovative Deep Neural Network Architectures for NLP with Python, PyTorch, TensorFlow, BERT, RoBERTa, and More

Creating and Deploying Deep Learning Applications

A practical approach to building neural network models using PyTorch

Become an AI language understanding expert by mastering the quantum leap of Transformer neural network models Key Features Build 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 models Go through hands-on applications in Python using Google Colaboratory Notebooks with nothing to install on a local machine Learn training tips and alternative language understanding methods to illustrate important key concepts Book 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 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 deep learning & NLP practitioners, data analysts and data scientists who want an introduction to AI language understanding to process the increasing amounts of language-driven functions.

One-stop solution for NLP practitioners, ML developers, and data scientists to build effective NLP systems that can perform real-world complicated tasks Key FeaturesApply deep learning algorithms and techniques such as BiLSTMS, CRFs, BPE and more using TensorFlow 2Explore applications like text generation, summarization, weakly supervised labelling and moreRead cutting edge material with seminal papers provided in the GitHub repository with full working codeBook Description Recently, there have been tremendous advances in NLP, and we are now moving from research labs into practical applications. This book comes with a perfect blend of both the theoretical and practical aspects of trending and complex NLP techniques. The book is focused on innovative applications in the field of NLP, language generation, and dialogue systems. It helps you apply the concepts of pre-processing text using techniques such as tokenization, parts of speech tagging, and lemmatization using popular libraries such as Stanford NLP and SpaCy. You will build Named Entity Recognition (NER) from scratch using Conditional Random Fields and Viterbi Decoding on top of RNNs. The book covers key emerging areas such as generating text for use in sentence completion and text summarization, bridging images and text by generating captions for images, and managing dialogue aspects of chatbots. You will learn how to apply transfer learning and fine-tuning using TensorFlow 2. Further, it covers practical techniques that can simplify the labelling of textual data. The book also has a working code that is adaptable to your use cases for each tech piece. By the end of the book, you will have an advanced knowledge of the tools, techniques and deep learning architecture used to solve complex NLP problems. What

you will learn Grasp important pre-steps in building NLP applications like POS tagging Use transfer and weakly supervised learning using libraries like Snorkel Do sentiment analysis using BERT Apply encoder-decoder NN architectures and beam search for summarizing texts Use Transformer models with attention to bring images and text together Build apps that generate captions and answer questions about images using custom Transformers Use advanced TensorFlow techniques like learning rate annealing, custom layers, and custom loss functions to build the latest DeepNLP models Who this book is for This is not an introductory book and assumes the reader is familiar with basics of NLP and has fundamental Python skills, as well as basic knowledge of machine learning and undergraduate-level calculus and linear algebra. The readers who can benefit the most from this book include intermediate ML developers who are familiar with the basics of supervised learning and deep learning techniques and professionals who already use TensorFlow/Python for purposes such as data science, ML, research, analysis, etc.

"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 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

Neural Network Methods in Natural Language Processing

Natural Language Processing Recipes

Build smart, AI-driven linguistic applications using deep learning and NLP techniques

PyTorch Recipes

Practical applications with deep learning

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

Deep Learning with PyTorch

Natural Language Processing with PyTorch Build Intelligent Language Applications Using Deep Learning"O'Reilly Media, Inc."

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 Learn how to use and implement transformers with Hugging Face and OpenAI (and others) by reading, running examples, investigating issues, asking the author questions, and interacting with our AI/ML community Key Features Pretrain a BERT-based model from scratch using Hugging Face Fine-tune powerful transformer models, including OpenAI's GPT-3, to learn the logic of your data Perform root cause analysis on hard NLP problems Book Description Transformers are...well...transforming the world of AI. There are many platforms and models out there, but which ones best suit

your needs? Transformers for Natural Language Processing, 2nd Edition, guides you through the world of transformers, highlighting the strengths of different models and platforms, while teaching you the problem-solving skills you need to tackle model weaknesses. You'll use Hugging Face to pretrain a RoBERTa model from scratch, from building the dataset to defining the data collator to training the model. If you're looking to fine-tune a pretrained model, including GPT-3, then Transformers for Natural Language Processing, 2nd Edition, shows you how with step-by-step guides. The book investigates machine translations, speech-to-text, text-to-speech, question-answering, and many more NLP tasks. It provides techniques to solve hard language problems and may even help with fake news anxiety (read chapter 13 for more details). You'll see how cutting-edge platforms, such as OpenAI, have taken transformers beyond language into computer vision tasks and code creation using Codex. By the end of this book, you'll know how transformers work and how to implement them and resolve issues like an AI detective! What you will learn Find out how ViT and CLIP label images (including blurry ones!) and create images from a sentence using DALL-E Discover new techniques to investigate complex language problems Compare and contrast the results of GPT-3 against T5, GPT-2, and BERT-based transformers Carry out sentiment analysis, text summarization, casual speech analysis, machine translations, and more using TensorFlow, PyTorch, and GPT-3 Measure the productivity of key transformers to define their scope, potential, and limits in production Who this book is for If you want to learn about and apply transformers to your natural language (and image) data, this book is for you. You'll need a good understanding of Python and deep learning and a basic understanding of NLP to benefit most from this book. Many platforms covered in this book provide interactive user interfaces, which allow readers with a general interest in NLP and AI to follow several chapters. And, don't worry if you get stuck or have questions; this book gives you direct access to our AI/ML community and author, Denis Rothman. So, he'll be there to guide you on your transformers journey!

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.

Linguistic Fundamentals for Natural Language Processing II Build, train, and fine-tune deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, and GPT-3

Natural Language Processing in Action

Creating Neural Networks with Python

100 Essentials from Semantics and Pragmatics

Natural Language Processing with TensorFlow

A Problem-Solution Approach

Learn to harness the power of AI for natural language processing, performing tasks such as spell check, text summarization, document classification, and natural language generation. Along the way, you will learn the skills to implement these methods in larger infrastructures to replace existing code or create new algorithms.

Applied Natural Language Processing with Python starts with reviewing the necessary machine learning concepts before moving onto discussing various NLP problems. After reading this book, you will have the skills to apply these concepts in your own professional environment. What You Will Learn Utilize various machine learning and natural language processing libraries such as TensorFlow, Keras, NLTK, and Gensim Manipulate and preprocess raw text data in formats such as .txt and .pdf Strengthen your skills in data science by learning both the theory and the application of various algorithms Who This Book Is For You should be at least a beginner in ML to get the most out of this text, but you needn't feel that you need be an expert to understand the content.

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 learnUnderstand classical linguistics in using English grammar for automatically generating questions and answers from a free text corpusWork with text embedding models for dense number representations of words, subwords and characters in the English language for exploring document clusteringDeep Learning in NLP using PyTorch with a code-driven introduction to PyTorchUsing an NLP project management Framework for estimating timelines and organizing your project into stagesHack and build a simple chatbot application in 30 minutesDeploy an NLP or machine learning application using Flask as RESTFUL APIsWho 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. Become a proficient NLP data scientist by developing deep learning models for NLP and extract valuable insights from structured and unstructured data Key FeaturesGet to grips with word embeddings, semantics, labeling, and high-level word representations using practical examplesLearn modern approaches to NLP and explore state-of-the-art NLP models using PyTorchImprove your NLP applications with innovative neural networks such as RNNs, LSTMs, and CNNsBook Description In the internet age, where an increasing volume of text data is generated daily from social media and other platforms, being able to make sense of that data is a crucial skill. With this book, you ' ll learn how to extract valuable insights from text by building deep learning models for natural language processing (NLP) tasks. Starting by understanding how to install PyTorch and using CUDA to accelerate the processing speed, you ' ll explore how the NLP architecture works with the help of practical examples. This PyTorch NLP book will guide you through core concepts such as word embeddings, CBOW, and tokenization in PyTorch. You ' ll then learn techniques for processing textual data and see how deep learning can be used for NLP tasks. The book demonstrates how to implement deep learning and neural network architectures to build models that will allow you to classify and translate text and perform sentiment analysis. Finally, you ' ll learn how to build advanced NLP models, such as conversational chatbots. By the end of this book, you ' ll not only have understood the different NLP problems that can be solved using deep learning with PyTorch, but also be able to build models to solve them. What you will learnUse NLP techniques for understanding, processing, and generating textUnderstand PyTorch, its applications and how it can be used to build deep linguistic modelsExplore the wide variety of deep learning architectures for NLPDevelop the skills you need to process and represent both structured and unstructured NLP dataBecome well-versed with state-of-the-art technologies and exciting new developments in the NLP domainCreate chatbots using attention-based neural networksWho this book is for This PyTorch book is for NLP developers, machine learning and deep learning developers, and anyone interested in building intelligent language applications using both traditional NLP approaches and deep learning architectures. If you ' re looking to adopt modern NLP techniques and models for your development projects, this book is for you. Working knowledge of Python programming, along with basic working knowledge of NLP tasks, is required. Build custom NLP models in record time by adapting pre-trained machine learning models to solve specialized problems. Summary In Transfer Learning for Natural Language Processing you will learn: Fine tuning pretrained models with new domain data Picking the right model to reduce resource usage Transfer learning for neural network architectures Generating text with generative pretrained transformers Cross-

lingual transfer learning with BERT Foundations for exploring NLP academic literature Training deep learning NLP models from scratch is costly, time-consuming, and requires massive amounts of data. In Transfer Learning for Natural Language Processing, DARPA researcher Paul Azunre reveals cutting-edge transfer learning techniques that apply customizable pretrained models to your own NLP architectures. You ' ll learn how to use transfer learning to deliver state-of-the-art results for language comprehension, even when working with limited label data. Best of all, you ' ll save on training time and computational costs. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Build custom NLP models in record time, even with limited datasets! Transfer learning is a machine learning technique for adapting pretrained machine learning models to solve specialized problems. This powerful approach has revolutionized natural language processing, driving improvements in machine translation, business analytics, and natural language generation. About the book Transfer Learning for Natural Language Processing teaches you to create powerful NLP solutions quickly by building on existing pretrained models. This instantly useful book provides crystal-clear explanations of the concepts you need to grok transfer learning along with hands-on examples so you can practice your new skills immediately. As you go, you ' ll apply state-of-the-art transfer learning methods to create a spam email classifier, a fact checker, and more real-world applications. What's inside Fine tuning pretrained models with new domain data Picking the right model to reduce resource use Transfer learning for neural network architectures Generating text with pretrained transformers About the reader For machine learning engineers and data scientists with some experience in NLP. About the author Paul Azunre holds a PhD in Computer Science from MIT and has served as a Principal Investigator on several DARPA research programs. Table of Contents PART 1 INTRODUCTION AND OVERVIEW 1 What is transfer learning? 2 Getting started with baselines: Data preprocessing 3 Getting started with baselines: Benchmarking and optimization PART 2 SHALLOW TRANSFER LEARNING AND DEEP TRANSFER LEARNING WITH RECURRENT NEURAL NETWORKS (RNNS) 4 Shallow transfer learning for NLP 5 Preprocessing data for recurrent neural network deep transfer learning experiments 6 Deep transfer learning for NLP with recurrent neural networks PART 3 DEEP TRANSFER LEARNING WITH TRANSFORMERS AND ADAPTATION STRATEGIES 7 Deep transfer learning for NLP with the transformer and GPT 8 Deep transfer learning for NLP with BERT and multilingual BERT 9 ULMFiT and knowledge distillation adaptation strategies 10 ALBERT, adapters, and multitask adaptation strategies 11 Conclusions

Getting started with Deep Learning for Natural Language Processing

Getting Started with Google BERT

Hands-On Natural Language Processing with Python

Hands-On Python Natural Language Processing

Applied Natural Language Processing with Python

Building Chatbots with Python

Mastering Transformers

Build neural network models in text, vision and advanced analytics using PyTorch Key Features Learn PyTorch for implementing cutting-edge deep learning algorithms. Train your neural networks for higher speed and flexibility and learn how to implement them in various scenarios; Cover various advanced neural network architecture such as ResNet, Inception, DenseNet and more with practical examples; Book Description Deep learning powers the most intelligent systems in the world, such as Google Voice, Siri, and Alexa. Advancements in powerful hardware, such as GPUs, software frameworks such as PyTorch, Keras, Tensorflow, and CNTK along with the availability of big data have made it easier to implement solutions to problems in the areas of text, vision, and advanced analytics. This book will get you up and running with one of the most cutting-edge deep learning libraries—PyTorch. PyTorch is grabbing the attention of deep learning researchers and data science professionals due to its accessibility, efficiency and

being more native to Python way of development. You'll start off by installing PyTorch, then quickly move on to learn various fundamental blocks that power modern deep learning. You will also learn how to use CNN, RNN, LSTM and other networks to solve real-world problems. This book explains the concepts of various state-of-the-art deep learning architectures, such as ResNet, DenseNet, Inception, and Seq2Seq, without diving deep into the math behind them. You will also learn about GPU computing during the course of the book. You will see how to train a model with PyTorch and dive into complex neural networks such as generative networks for producing text and images. By the end of the book, you'll be able to implement deep learning applications in PyTorch with ease. What you will learn

- Use PyTorch for GPU-accelerated tensor computations
- Build custom datasets and data loaders for images and test the models using torchvision and torchtext
- Build an image classifier by implementing CNN architectures using PyTorch
- Build systems that do text classification and language modeling using RNN, LSTM, and GRU
- Learn advanced CNN architectures such as ResNet, Inception, Densenet, and learn how to use them for transfer learning
- Learn how to mix multiple models for a powerful ensemble model
- Generate new images using GAN 's and generate artistic images using style transfer

Who this book is for This book is for machine learning engineers, data analysts, data scientists interested in deep learning and are looking to explore implementing advanced algorithms in PyTorch. Some knowledge of machine learning is helpful but not a mandatory need. Working knowledge of Python programming is expected.

"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.

Implement natural language processing applications with Python using a problem-solution approach. This book has numerous coding exercises that will help you to quickly deploy natural language processing techniques, such as text classification, parts of speech identification, topic modeling, text summarization, text generation, entity extraction, and sentiment analysis. Natural Language Processing Recipes starts by offering solutions for cleaning and preprocessing text data and ways to analyze it with advanced algorithms. You ' ll see practical applications of the semantic as well as syntactic analysis of text, as well as complex natural language processing approaches that involve text normalization, advanced preprocessing, POS tagging, and sentiment analysis. You will also learn various applications of machine learning and deep learning in natural language processing. By using the recipes in this book, you will have a toolbox of solutions to apply to your own projects in the real world, making your development time quicker and more efficient. What You Will LearnApply NLP techniques using Python libraries such as NLTK, TextBlob, spaCy, Stanford CoreNLP, and many more Implement the concepts of information retrieval, text summarization, sentiment analysis, and other advanced natural language processing techniques. Identify machine learning and deep learning techniques for natural language processing and natural language generation problems Who This Book Is ForData scientists who want to refresh and learn various concepts of natural language processing through coding exercises.

Write modern natural language processing applications using deep learning algorithms and TensorFlow Key Features Focuses on more efficient natural language processing using TensorFlow Covers NLP as a field in its own right to improve understanding for choosing TensorFlow tools and other deep learning approaches Provides choices for how to process and evaluate large unstructured text datasets Learn to apply the TensorFlow toolbox to specific tasks in the most interesting field in artificial intelligence Book Description Natural language processing (NLP) supplies the majority of data available to deep learning applications, while TensorFlow is the most important deep learning framework currently available. Natural Language Processing with TensorFlow brings TensorFlow and NLP together to give you invaluable tools to work with the immense volume of unstructured data in today ' s data streams, and apply these tools to specific NLP tasks. Thushan Ganegedara starts by giving you a grounding in NLP and TensorFlow basics. You'll then learn how to use Word2vec, including advanced extensions, to create word embeddings that turn sequences of words into vectors accessible to deep learning algorithms. Chapters on classical deep learning algorithms, like convolutional neural networks (CNN) and recurrent neural networks (RNN), demonstrate important NLP tasks as sentence classification and language generation. You will learn how to apply high-performance RNN models, like long short-term memory (LSTM) cells, to NLP tasks. You will also explore neural machine translation and

implement a neural machine translator. After reading this book, you will gain an understanding of NLP and you'll have the skills to apply TensorFlow in deep learning NLP applications, and how to perform specific NLP tasks. What you will learn Core concepts of NLP and various approaches to natural language processing How to solve NLP tasks by applying TensorFlow functions to create neural networks Strategies to process large amounts of data into word representations that can be used by deep learning applications Techniques for performing sentence classification and language generation using CNNs and RNNs About employing state-of-the-art advanced RNNs, like long short-term memory, to solve complex text generation tasks How to write automatic translation programs and implement an actual neural machine translator from scratch The trends and innovations that are paving the future in NLP Who this book is for This book is for Python developers with a strong interest in deep learning, who want to learn how to leverage TensorFlow to simplify NLP tasks. Fundamental Python skills are assumed, as well as some knowledge of machine learning and undergraduate-level calculus and linear algebra. No previous natural language processing experience required, although some background in NLP or computational linguistics will be helpful.

Natural Language Processing with Transformers

A practical guide to applying deep learning architectures to your NLP applications

Natural Language Processing with PyTorch

Applied Natural Language Processing in the Enterprise

Analyzing Text with the Natural Language Toolkit

Speech & Language Processing

Get up to speed with the deep learning concepts of Pytorch using a problem-solution approach. Starting with an introduction to PyTorch, you'll get familiarized with tensors, a type of data structure used to calculate arithmetic operations and also learn how they operate. You will then take a look at probability distributions using PyTorch and get acquainted with its concepts. Further you will dive into transformations and graph computations with PyTorch. Along the way you will take a look at common issues faced with neural network implementation and tensor differentiation, and get the best solutions for them. Moving on to algorithms; you will learn how PyTorch works with supervised and unsupervised algorithms. You will see how convolutional neural networks, deep neural networks, and recurrent neural networks work using PyTorch. In conclusion you will get acquainted with natural language processing and text processing using PyTorch. What You Will Learn Master tensor operations for dynamic graph-based calculations using PyTorch Create PyTorch transformations and graph computations for neural networks Carry out supervised and unsupervised learning using PyTorch Work with deep learning algorithms such as CNN and RNN Build LSTM models in PyTorch Use PyTorch for text processing Who This Book Is For Readers wanting to dive straight into programming PyTorch.

Foster your NLP applications with the help of deep learning, NLTK, and TensorFlow

Key Features

- Weave neural networks into linguistic applications across various platforms**
- Perform NLP tasks and train its models using NLTK and TensorFlow**
- Boost your NLP models with strong deep learning architectures such as CNNs and RNNs**

Book Description

Natural language processing (NLP) has found its application in various domains, such as web search, advertisements, and customer services, and with the help of deep learning, we can enhance its performances in these areas. *Hands-On Natural Language Processing with Python* teaches you how to leverage deep learning models for performing various NLP tasks, along with best practices in dealing with today's NLP challenges. To begin with, you will understand the core concepts of NLP and deep learning, such as Convolutional Neural Networks (CNNs), recurrent neural networks (RNNs), semantic embedding, Word2vec, and more. You will learn how to perform each and every task of NLP using neural networks, in which you will train and deploy neural networks in your NLP applications. You will get accustomed to using RNNs and CNNs in various application areas, such as text classification and sequence labeling, which are essential in the application of sentiment analysis, customer service chatbots, and anomaly detection. You will be equipped with practical knowledge in order to implement deep learning in your linguistic applications using Python's popular deep learning library, TensorFlow. By the end of this book, you will be well versed in building deep learning-backed NLP applications, along with overcoming NLP challenges with best practices developed by domain experts.

What you will learn

- Implement semantic embedding of words to classify and find entities**
- Convert words to vectors by training in order to perform arithmetic operations**
- Train a deep learning model to detect classification of tweets and news**
- Implement a question-answer model with search and RNN models**
- Train models for various text classification datasets using CNN**
- Implement WaveNet a deep generative model for producing a natural-sounding voice**
- Convert voice-to-text and text-to-voice**
- Train a model to convert speech-to-text using DeepSpeech**

Who this book is for

Hands-on Natural Language Processing with Python is for you if you are a developer, machine learning or an NLP engineer who wants to build a deep learning application that leverages NLP techniques. This comprehensive guide is also useful for deep learning users who want to extend their deep learning skills in building NLP applications. All you need is the basics of machine learning and Python to enjoy the book.

Humans do a great job of reading text, identifying key ideas, summarizing, making connections, and other tasks that require comprehension and context. Recent advances in deep learning make it possible for computer systems to achieve similar results. *Deep Learning for Natural Language Processing* teaches you to apply deep learning methods to natural

language processing (NLP) to interpret and use text effectively. In this insightful book, NLP expert Stephan Raaijmakers distills his extensive knowledge of the latest state-of-the-art developments in this rapidly emerging field. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Build your own chatbot using Python and open source tools. This book begins with an introduction to chatbots where you will gain vital information on their architecture. You will then dive straight into natural language processing with the natural language toolkit (NLTK) for building a custom language processing platform for your chatbot. With this foundation, you will take a look at different natural language processing techniques so that you can choose the right one for you. The next stage is to learn to build a chatbot using the API.ai platform and define its intents and entities. During this example, you will learn to enable communication with your bot and also take a look at key points of its integration and deployment. The final chapter of *Building Chatbots with Python* teaches you how to build, train, and deploy your very own chatbot. Using open source libraries and machine learning techniques you will learn to predict conditions for your bot and develop a conversational agent as a web application. Finally you will deploy your chatbot on your own server with AWS. **What You Will Learn** Gain the basics of natural language processing using Python Collect data and train your data for the chatbot Build your chatbot from scratch as a web app Integrate your chatbots with Facebook, Slack, and Telegram Deploy chatbots on your own server **Who This Book Is For** Intermediate Python developers who have no idea about chatbots. Developers with basic Python programming knowledge can also take advantage of the book.

Teach language to machines using Python's deep learning library

Deep Learning in Natural Language Processing

Implementing Machine Learning and Deep Learning Algorithms for Natural Language Processing

Advanced Natural Language Processing with TensorFlow 2

Transformers for Natural Language Processing

Deep Learning for Natural Language Processing

Introduction to Natural Language Processing

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.

Getting Started with Google BERT will help you become well-versed with the BERT model from scratch and learn how to create interesting NLP applications. You'll understand several variants of BERT such as ALBERT, RoBERTa, DistilBERT, ELECTRA, VideoBERT, and many others in detail.

Take a problem-solving approach to learning all about transformers and get up and running in no time by implementing methodologies that will build the future of NLP

Key Features

- Explore quick prototyping with up-to-date Python libraries to create effective solutions to industrial problems
- Solve advanced NLP problems such as named-entity recognition, information extraction, language generation, and conversational AI
- Monitor your model's performance with the help of BertViz, exBERT, and TensorBoard

Book Description Transformer-based language models have dominated natural language processing (NLP) studies and have now become a new paradigm. With this book, you'll learn how to build various transformer-based NLP applications using the Python Transformers library. The book gives you an introduction to Transformers by showing you how to write your first hello-world program. You'll then learn how a tokenizer works and how to train your own tokenizer. As you advance, you'll explore the architecture of autoencoding models, such as BERT, and autoregressive models, such as GPT. You'll see how to train and fine-tune models for a variety of natural language understanding (NLU) and natural language generation (NLG) problems, including text classification, token classification, and text representation. This book also helps you to learn efficient models for challenging problems, such as long-context NLP tasks with limited computational capacity. You'll also work with multilingual and cross-lingual problems, optimize models by monitoring their performance, and discover how to deconstruct these models for interpretability and explainability. Finally, you'll be able to deploy your transformer models in a production environment. By the end of this NLP book, you'll have learned how to use Transformers to solve advanced NLP problems using advanced models. What you will learn

- Explore state-of-the-art NLP solutions with the Transformers library
- Train a language model in any language with any transformer architecture
- Fine-tune a pre-trained language model to

perform several downstream tasks • Select the right framework for the training, evaluation, and production of an end-to-end solution • Get hands-on experience in using TensorBoard and Weights & Biases • Visualize the internal representation of transformer models for interpretability

Who this book is for This book is for deep learning researchers, hands-on NLP practitioners, as well as ML/NLP educators and students who want to start their journey with Transformers. Beginner-level machine learning knowledge and a good command of Python will help you get the best out of this book.

Table of Contents • From Bag-of-Words to the Transformers • A Hands-On Introduction to the Subject • Autoencoding Language Models • Autoregressive and Other Language Models • Fine-Tuning Language Models for Text Classification • Fine-Tuning Language Models for Token Classification • Text Representation • Working with Efficient Transformers • Cross-Lingual and Multilingual Language Modeling • Serving Transformer Models • Attention Visualization and Experiment Tracking

Review "Transformers rule for a lot of NLP tasks now, and this is a great book about them. Beginners will appreciate clear explanations and experienced programmers have plenty of examples how to use Transformers even for complex tasks. Code examples are well selected and I did like that they use both Tensorflow and PyTorch." -- Andrzej Jankowski, AI Sales Engineer at Intel and Business AI Postgraduate Course Leader at Kozminski University

Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. In Real-world Natural Language Processing you will learn how to: Design, develop, and deploy useful NLP applications

Create named entity taggers

Build machine translation systems

Construct language generation systems and chatbots

Use advanced NLP concepts such as attention and transfer learning

Real-world Natural Language Processing teaches you how to create practical NLP applications without getting bogged down in complex language theory and the mathematics of deep learning. In this engaging book, you'll explore the core tools and techniques required to build a huge range of powerful NLP apps, including chatbots, language detectors, and text classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About the technology

Training computers to interpret and generate speech and text is a monumental challenge, and the payoff for reducing labor and improving human/computer interaction is huge! The field of Natural Language Processing (NLP) is advancing rapidly, with countless new tools and practices. This unique book offers an innovative collection of NLP techniques with applications in machine translation, voice assistants, text generation, and more.

About the book

Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. Guided by clear explanations of each core NLP topic, you'll create many interesting applications including a sentiment analyzer and a chatbot. Along the way, you'll use Python and open source libraries like AllenNLP and HuggingFace Transformers to speed up your development process. What's inside

Design, develop, and deploy useful NLP applications

Create named entity taggers

Build machine translation systems

Construct language generation systems and chatbots

About the reader

For Python programmers. No prior machine learning knowledge assumed.

About the author

Masato Hagiwara received his computer science PhD from Nagoya University in 2009. He has interned at Google and Microsoft Research, and worked at Duolingo as a Senior Machine Learning Engineer. He now runs his own research and consulting company.

Table of Contents

PART 1 BASICS

1 Introduction to natural language processing

2 Your first NLP application

3 Word and document embeddings

4 Sentence classification

5 Sequential labeling and language modeling

PART 2 ADVANCED MODELS

6 Sequence-to-sequence models

7 Convolutional neural networks

8 Attention and Transformer

9 Transfer learning with pretrained language models

PART 3 PUTTING INTO PRODUCTION

10 Best practices in developing NLP applications

11 Deploying and serving NLP applications

Hands-On Natural Language Processing with PyTorch 1.x

Natural Language Processing with Python

Deep Learning Illustrated

Real-World Natural Language Processing

A Comprehensive Guide to Building Real-World NLP Systems

Natural Language Processing with Python Quick Start Guide

Explore deep learning concepts and implement over 50 real-world image applications

Meaning is a fundamental concept in Natural Language Processing (NLP), in the tasks of both Natural Language Understanding (NLU) and Natural Language Generation (NLG). This is because the aims of these fields are to build systems that understand what people mean when they speak or write, and that can produce linguistic strings that successfully express to people the intended content. In order for NLP to scale beyond partial, task-specific solutions, researchers in these fields must be informed by what is known about how humans use language to express and understand communicative intents. The purpose of this book is to present a selection of useful information about semantics and pragmatics, as understood in linguistics, in a way that's accessible to and useful for NLP practitioners with minimal (or even no) prior training in linguistics.

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)

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.

Deep Learning for Coders with fastai and PyTorch