

Cs224d Tensorflow Tutorial

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

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

Guaranteeing you 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 Poster 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

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.

Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques
A beginner's guide to getting up and running with deep learning from scratch using Python
A Visual Approach

Apply modern RL methods, with deep Q-networks, value iteration, policy gradients, TRPO, AlphaGo Zero and more
Develop Deep Learning Models for your Natural Language Problems
Deep Learning in Natural Language Processing

Innovations in Computer Science and Engineering
Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide About This Book Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide Real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is intended for a general audience of people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. What You Will Learn Learn about machine learning landscapes along with the historical development and progress of deep learning Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x Access public datasets and utilize them using TensorFlow to load, process, and transform data Use TensorFlow on real-world datasets, including images, text, and more Learn how to evaluate the performance of your deep learning models Using deep learning for scalable object detection and mobile computing Train machines quickly to learn from data by exploring reinforcement learning techniques Explore active areas of deep learning research and applications In Detail Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples, you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing.

On the surface, research on machine learning and data mining has led to a wide variety of algorithms that learn general functions from experience. As machine learning is maturing, it has begun to make the successful transition from academic research to various practical applications. Generic techniques such as decision trees and artificial neural networks, for example, are now being used in various commercial and industrial applications. Learning to Learn is an exciting new research direction within machine learning. Similar to traditional machine-learning algorithms, the methods described in Learning to Learn induce general functions from experience. However, the book investigates algorithms that can change the way they generalize, i.e., practice the task of learning itself, and improve on it. To illustrate the utility of learning to learn, it is worthwhile comparing machine learning with human learning. Humans encounter a continual stream of learning tasks. They do not just learn concepts or motor skills, they also learn bias, i.e., they learn how to generalize. As a result, humans are often able to generalize correctly from extremely few examples - often just a single example suffices to teach us a new thing. A deeper understanding of computer programs that improve their ability to learn can have a large practical impact on the field of machine learning and beyond. In recent years, the field has made significant progress towards a theory of learning to learn along with practical new algorithms, some of which led to impressive results in real-world applications. Learning to Learn provides a survey of some of the most exciting new research concepts, written by leading researchers in the field. Its objective is to investigate the utility and feasibility of computer programs that can learn how to learn, both from a practical and a theoretical point of view.

This book gathers a collection of high-quality peer-reviewed research papers presented at the 2nd International Conference on Data and Information Sciences (ICDIS 2019), held at Raja Balwant Singh Engineering Technical Campus, Agra, India, on March 29 – 30, 2019. In chapters written by leading researchers, developers, and practitioner from academia and industry, it covers virtually all aspects of computational sciences and information security, including central topics like artificial intelligence, cloud computing, and big data. Highlighting the latest developments and technical solutions, it will show readers from the computer industry how to capitalize on key advances in next-generation computer and communication technology.

Learn how to apply TensorFlow to a wide range of deep learning and Machine Learning problems with this practical guide on training CNNs for image classification, image recognition, object detection and many computer vision challenges. Key Features Learn the fundamentals of Convolutional Neural Networks Harness Python and TensorFlow to train CNNs Build scalable deep learning models that can process millions of items Book Description Convolutional Neural Networks (CNN) are one of the most popular architectures used in computer vision apps. This book is an introduction to CNNs through solving real-world problems in deep learning while teaching you their implementation in popular Python library - TensorFlow. By the end of the book, you will be training CNNs in no time! We start with an overview of popular machine learning and deep learning models, and then get you set up with a TensorFlow development environment. This environment is the basis for implementing and training deep learning models in later chapters. Then, you will use Convolutional Neural Networks to work on problems such as image classification, object detection, and semantic segmentation. After that, you will use transfer learning to see how these models can solve other deep learning problems. You will also get a taste of implementing generative models such as autoencoders and generative adversarial networks. Later on, you will use useful tips on machine learning best practices and troubleshooting. Finally, you will learn how to apply your models on large datasets of millions of images. What you will learn Train machine learning models with TensorFlow Create systems that can evolve and scale during their life cycle Use CNNs in image recognition and classification Use TensorFlow for building deep learning models Train popular deep learning models Fine-tune a neural network to improve the quality of results with transfer learning Build TensorFlow models that can scale to large datasets and systems Who this book is for This book is for Software Engineers, Data Scientists, or Machine Learning practitioners who want to use CNNs for solving real-world problems. Knowledge of basic machine learning concepts, linear algebra and Python will help.

Teaching Machines to Paint, Write, Compose, and Play
Develop Deep Learning Models on Theano and TensorFlow Using Keras

Natural Language Processing with PyTorch
A Guide to Data Science for Fraud Detection

Programming PyTorch for Deep Learning
Build Intelligent Language Applications Using Deep Learning

The Master Algorithm
Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features! Learn and implement machine learning algorithms in a variety of real-life scenariosCover a range of tasks catering to supervised, unsupervised and reinforcement learning techniquesFind easy-to-follow code solutions for tackling common and not-so-common challengesBook Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learnUse predictive modeling and apply it to real-world problemsExplore data visualization techniques to interact with your dataLearn how to build a recommendation engineUnderstand how to interact with text data and build models to analyze itWork with speech data and recognize spoken words using Hidden Markov ModelsGet well versed with reinforcement learning, automated ML, and transfer learningWork with image data and build systems for image recognition and biometric face recognitionUse deep neural networks to build an optical character recognition systemWho this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need.

Familiarity with Python programming and machine learning concepts will be useful. You'll learn how to use key deep learning algorithms without the need for complex math. Ever since computers began beating us at chess, they've been getting better at a wide range of human activities, from writing songs and generating news articles to helping doctors provide healthcare. Deep learning is the source of many of these breakthroughs, and its remarkable ability to find patterns hiding in data has made it the fastest growing field in artificial intelligence (AI). Digital assistants on our phones use deep learning to understand and respond intelligently to voice commands; automotive systems use it to safely navigate road hazards; online platforms use it to deliver personalized suggestions for movies and books - the possibilities are endless. Deep Learning: A Visual Approach is for anyone who wants to understand this fascinating field and programming usually required to grasp its internals. If you want to know how these tools work, and use them yourself, the answers are all within these pages. And, if you're ready to write your own programs, there are also plenty of supplemental Python notebooks in the accompanying Github repository to get you going. The book's conversational style, extensive color illustrations, illuminating analogies, and real-world examples expertly explain the key concepts in deep learning, including: • How text generators create novel stories and articles • How deep learning systems learn to play and win at human games • How image classification systems identify objects or people in a photo • How to think about probabilities in a way that's useful to everyday life • How to use the machine learning techniques that form the core of modern AI

Intellectual adventures of all kinds can use the powerful ideas covered in Deep Learning: A Visual Approach to build intelligent systems that help us better understand the world and everyone who lives in it. It's the future of AI, and this book allows you to fully envision it. Full Color Illustrations
Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more Key FeaturesDeep Learning solutions from Kaggle Masters and Google Developer ExpertsGet to grips with the fundamentals including variables, matrices, and data sourcesLearn advanced techniques to make your algorithms faster and more accurateBook Description The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning - each using Google's machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources. You'll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learnTake TensorFlow into productionImplement and fine-tune Transformer models for various NLP tasksApply reinforcement learning algorithms using the TF-Agents frameworkUnderstand linear regression techniques and use Estimators to train linear modelsExecute neural networks and improve predictions on tabular dataMaster convolutional neural networks and recurrent neural networks through practical recipesWho this book is for If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

Sentiment analysis is a branch of natural language processing concerned with the study of the intensity of the emotions expressed in a piece of text. The automated analysis of the multitude of messages delivered through social media is one of the hottest research fields, both in academy and in industry, due to its extremely high potential applicability in many different domains. This Special Issue describes both technological contributions to the field, mostly based on deep learning techniques, and specific applications in areas like health insurance, gender classification, recommender systems, and cyber aggression detection.
Deep Learning for Coders with fastai and PyTorch
Proceedings of 5th International Conference on the Industry 4.0 Model for Advanced Manufacturing
A Hands-on Introduction to Learning Algorithms
Introduction to Deep Learning
Graph Representation Learning
Generative Deep Learning

Solve computer vision problems with modeling in TensorFlow and Python
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

This book includes research-based approaches for sentiment analysis, a relatively new, but fast-growing research area, which has significantly changed in the past few years. The book presents a collection of state-of-the-art approaches, focusing on the best performing, cutting-edge solutions for the most common and difficult challenges faced in sentiment analysis research. Providing detailed explanations of the methodologies, the book is a valuable resource for researchers as well as newcomers to the field.

The core of this paper is a general set of variational principles for the problems of computing marginal probabilities and modes, applicable to multivariate statistical models in the exponential family. This book includes high-quality, peer-reviewed research papers from the 6thInternational Conference on Innovations in Computer Science & Engineering (ICICSE 2018), held at Guru Nanak Institutions, Hyderabad, India from August 17 to 18, 2018. The book discusses a wide variety of industrial, engineering and scientific applications of the emerging techniques and offers a platform for researchers from academia and industry to present their original work and exchange ideas, information, techniques and applications in the field of computer science.

A Mathematical Approach
Models, Learning, and Inference
Deep Learning Illustrated

Deep Learning
Deep Learning with PyTorch
Building Natural Language Generation Systems

Graphical Models, Exponential Families, and Variational Inference
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 basis 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

The book includes research papers from the 6thInternational Conference on Soft Computing and Signal Processing (ICSSP 2018). It features papers on current topics, such as soft sets, rough sets, fuzzy logic, neural networks, genetic algorithms and machine learning. It also discusses various aspects of these topics, like technologies, product implementation, and application issues.

Supervised sequence labelling is a vital area of machine learning, encompassing tasks such as speech, handwriting and gesture recognition, protein secondary structure prediction and part-of-speech tagging. Recurrent neural networks are powerful sequence learning tools—robust to input noise and distortion, able to exploit long-range contextual information—that were seen ideally suited to such problems. However their role in large-scale sequence labelling systems has so far been auxiliary. The goal of this book is a complete framework for classifying and transcribing sequential data with recurrent neural networks only. Three main innovations are introduced in order to realise this goal. Firstly, the connectionist temporal classification output layer allows the framework to be trained with unsegmented target sequences, such as phoneme-level speech transcriptions; this is in contrast to previous connectionist approaches, which were dependent on error-prone prior segmentation. Secondly, a multidimensional recurrent neural networks extend the framework in a natural way to data with more than one spatio-temporal dimension, such as images and videos. Thirdly, the use of hierarchical subsampling makes it feasible to apply the framework to very large or high resolution sequences, such as raw audio or video. Experimental validation is provided by state-of-the-art results in speech and handwriting recognition.

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep learning in particular—understand how a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source DeepLearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll see RL methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in neural networks Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec. DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

Neural Network Methods in Natural Language Processing
Python Machine Learning Cookbook
Deep Learning for Beginners
Proceedings of ICDIS 2019

Sentiment Analysis for Social Media
TensorFlow Machine Learning Cookbook
Deep Learning for Natural Language Processing

Implement supervised, unsupervised, and generative deep learning (DL) models using Keras and Dopamine with TensorFlow Key FeaturesUnderstand the fundamental machine learning concepts useful in deep learningLearn the underlying mathematical concepts as you implement deep learning models from scratchExplore easy-to-understand examples and use cases that will help you build a solid foundation in DLBook Description With information on the web exponentially increasing, it has become more difficult than ever to navigate through everything to find reliable content that will help you get started with deep learning. This book is designed to help you if you're a beginner looking to work on deep learning and build deep learning models from scratch, and you already have the basic mathematical and programming knowledge required to get started. The book begins with a basic overview of machine learning, guiding you through setting up popular Python frameworks. You will also understand how to prepare data by cleaning and preprocessing it for deep learning, and gradually go on to explore neural networks. A dedicated section will give you insights into the working of neural networks by helping you get hands-on with training single and multiple layers of neurons. Later, you will cover popular neural network architectures such as CNNs, RNNs, AEs, VAEs, and GANs with the help of simple examples, and learn how to build models from scratch. At the end of each chapter, you will find a question and answer section to help you test what you've learned through the course of the book. By the end of this book, you'll be well-versed with deep learning concepts and have the knowledge you need to use specific algorithms with various tools for different tasks. What you will learnImplement recurrent neural networks (RNNs) and long short-term memory (LSTM) for image classification and natural language processing tasksExplore the role of convolutional neural networks (CNNs) in computer vision and signal processingDiscover the ethical implications of deep learning modelingUnderstand the mathematical terminology associated with deep learningCode a generative adversarial network (GAN) and a variational autoencoder (VAE) to generate images from a learned latent spaceImplement visualization techniques to compare AEs and VAEsWho this book is for This book is for aspiring data scientists and deep learning engineers who want to get started with the fundamentals of deep learning and neural networks. Although no prior knowledge of deep learning or machine learning is required, familiarity with linear algebra and Python programming is necessary to get started.

" 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 specialty trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introduction: 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 nodes 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodes 14 End-to-end node analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Detect fraud earlier to mitigate loss and prevent cascading damage Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques is an authoritative guidebook for setting up a comprehensive fraud detection analytics solution. Early detection is a key factor in mitigating fraud damage, but it involves more specialized techniques than detecting fraud at the more advanced stages. This invaluable guide details both the theory and technical aspects of these techniques, and provides expert insight into streamlining implementation. Coverage includes data gathering, preprocessing, model building, and post-implementation, with comprehensive guidance on various learning techniques and the data types utilized by each. These techniques are effective for fraud detection across industry boundaries, including applications in insurance fraud, credit card fraud, anti-money laundering, healthcare fraud, telecommunications fraud, click fraud, tax evasion, and more, giving you a highly practical framework for fraud prevention. It is estimated that a typical organization loses about 5% of its revenue to fraud every year. More effective fraud detection is possible, and this book describes the various analytical techniques your organization must implement to put a stop to the revenue leak. Examine fraud patterns in historical data Utilize labeled, unlabeled, and networked data Detect fraud before the damage cascades Reduce losses, increase recovery, and tighten security The longer fraud is allowed to go on, the more harm it causes. It expands exponentially, sending ripples of damage throughout the organization, and becomes more and more complex to track, stop, and reverse. Fraud prevention relies on early and effective fraud detection, enabled by the techniques discussed here. Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques helps you stop fraud in its tracks, and eliminate the opportunities for future occurrence.

A modern treatment focusing on learning and inference, with minimal prerequisites, real-world examples and implementable algorithms.
TensorFlow for Machine Intelligence
A Practitioner's Approach

Supervised Sequence Labelling with Recurrent Neural Networks
Machine Learning in Action

Creating and Deploying Deep Learning Applications
Deep Learning Architectures

TensorFlow for Deep Learning
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 undergraduates and graduate students, post-doctoral researchers, lecturers and industrial researchers, as well as anyone interested in deep learning and natural language processing.

Dependency-based methods for syntactic parsing have become increasingly popular in natural language processing in recent years. This book gives a thorough introduction to the methods that are most widely used today. After an introduction to dependency grammar and dependency parsing, followed by a formal characterization of the dependency parsing problem, the book surveys the three major classes of parsing models that are in current use: transition-based, graph-based, and grammar-based models. It continues with a chapter on evaluation and one on the comparison of different methods, and it closes with a few words on current trends and future prospects of dependency parsing. The book presupposes a knowledge of basic concepts in linguistics and computer science, as well as some knowledge of parsing methods for constituency-based representations. List of Contents: Introduction / Dependency Parsing / Transition-Based Parsing / Graph-Based Parsing / Grammar-Based Parsing / Evaluation / Comparison / Final Thoughts

This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. Key Features Explore deep reinforcement learning (RL), from the first principles to the latest algorithms Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms Keep up with the very latest industry developments, including AI-driven chatbots Book Description Recent developments in reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning Hands-On is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks.

Deep learning is driving a revolution in AI, and it's only getting started. Discover how to deal with discrete and continuous action spaces in various environments Defeat Atari arcade games using the value iteration method Create your own OpenAI Gym environment to train a stock trading agent Teach your agent to play Connect4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI-driven chatbots Who this book is for Those who are fluent in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL.

"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 input from Grant Beyelved and beautiful illustrations by Agla & 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 nimble 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 eager 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.

Machine Learning Using TensorFlow Cookbook
Over 100 recipes to progress from smart data analytics to deep learning using real-world datasets, 2nd Edition
Hands-On Convolutional Neural Networks with TensorFlow
From Linear Regression to Reinforcement Learning

AMP 2020
Speech & Language Processing
Deep Learning With Python

Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced techniques that bring more accuracy and speed to machine learning Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning activities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before. You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, clustering analysis, artificial neural networks, and deep learning — each using Google's machine learning library TensorFlow. This guide starts with the fundamentals of the TensorFlow library which includes variables, matrices, and various data sources. Moving ahead, you will get hands-on experience with Linear Regression techniques with TensorFlow. The next chapters cover important high-level concepts such as neural networks, CNN, RNN, and NLP. Once you are familiar and comfortable with the TensorFlow ecosystem, the last chapter will show you how to take it to production. Style and approach This book takes a recipe-based approach where every topic is explicated with the help of a real-world example.

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 to 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

This book explains how to build Natural Language Generation (NLG) systems - computer software systems which use techniques from artificial intelligence and computational linguistics to automatically generate understandable texts in English or other human languages, either in isolation or as part of multimedia documents, Web pages, and speech output systems. Typically starting from some non-linguistic representation of information as input, NLG systems use knowledge about language and the application domain to automatically produce documents, reports, explanations, help messages, and other kinds of texts. The book covers the algorithms and representations needed to perform the core tasks of document planning, microplanning, and surface realization, using a case study to show how these components fit together. It also discusses engineering issues such as system architecture, requirements analysis, and the integration of text generation into multimedia and speech output systems.

Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.
Create powerful machine learning algorithms with TensorFlow
Deep Learning-Based Approaches for Sentiment Analysis
Dependency Parsing

Computer Vision

Proceedings of the Sixth IJCVSE 2018

Deep Reinforcement Learning Hands-On

Proceedings of ICSCSP 2018, Volume 2

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning.

This book gathers the proceedings of the 5th International Conference on the Industry 4.0 Model for Advanced Manufacturing (AMP 2020), held in Belgrade, Serbia, on 1 – 4 June 2020. The event marks the latest in a series of high-level conferences that bring together experts from academia and industry to exchange knowledge, ideas, experiences, research findings, and information in the field of manufacturing. The book addresses a wide range of topics, including: design of smart and intelligent products, developments in CAD/CAM technologies, rapid prototyping and reverse engineering, multistage manufacturing processes, manufacturing automation in the Industry 4.0 model, cloud-based products, and cyber-physical and reconfigurable manufacturing systems. By providing updates on key issues and highlighting recent advances in manufacturing engineering and technologies, the book supports the transfer of vital knowledge to the next generation of academics and practitioners. Further, it will appeal to anyone working or conducting research in this rapidly evolving field.

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-

Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

A Visual, Interactive Guide to Artificial Intelligence

Learning to Learn

Advances in Data and Information Sciences

Soft Computing and Signal Processing

Deep Learning with TensorFlow

How the Quest for the Ultimate Learning Machine Will Remake Our World

This book describes how neural networks operate from the mathematical point of view. As a result, neural networks can be interpreted both as function universal approximators and information processors. The book bridges the gap between ideas and concepts of neural networks, which are used nowadays at an intuitive level, and the precise modern mathematical language, presenting the best practices of the former and enjoying the robustness and elegance of the latter. This book can be used in a graduate course in deep learning, with the first few parts being accessible to senior undergraduates. In addition, the book will be of wide interest to machine learning researchers who are interested in a theoretical understanding of the subject.

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks

Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units