

Deep Learning Avec Tensorflow Mise En Oeuvre Et Cas Concrets

The official book on the Rust programming language, written by the Rust development team at the Mozilla Foundation, fully updated for Rust 2018. The Rust Programming Language is the official book on Rust: an open source systems programming language that helps you write faster, more reliable software. Rust offers control over low-level details (such as memory usage) in combination with high-level ergonomics, eliminating the hassle traditionally associated with low-level languages. The authors of The Rust Programming Language, members of the Rust Core Team, share their knowledge and experience to show you how to take full advantage of Rust's features--from installation to creating robust and scalable programs. You'll begin with basics like creating functions, choosing data types, and binding variables and then move on to more advanced concepts, such as:

- Ownership and borrowing, lifetimes, and traits
- Using Rust's memory safety guarantees to build fast, safe programs
- Testing, error handling, and effective refactoring
- Generics, smart pointers, multithreading, trait objects, and advanced pattern matching
- Using Cargo, Rust's built-in package manager, to build, test, and document your code and manage dependencies
- How best to use Rust's advanced compiler with compiler-led programming techniques

You'll find plenty of code examples throughout the book, as well as three chapters dedicated to building complete projects to test your learning: a number guessing game, a Rust implementation of a command line tool, and a multithreaded server. New to this edition: An extended section on Rust macros, an expanded chapter on modules, and appendixes on Rust development tools and editions.

This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for deep learning, Natural Language Processing (NLP), speech recognition, and general predictive analytics. The book provides a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. The authors begin by working through some basic examples in TensorFlow before diving deeper into topics such as CNN, RNN, LSTM, and GNN. The book is written for those who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. The authors demonstrate TensorFlow projects on Single Board Computers (SBCs).

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

Written by a pioneer of game development in academia, *Introduction to Game Physics with Box2D* covers the theory and practice of 2D game physics in a relaxed and entertaining yet instructional style. It offers a cohesive treatment of the topics and code involved in programming the physics for 2D video games. Focusing on writing elementary game physics code, the first half of the book helps you grasp the challenges of programming game physics from scratch, without libraries or outside help. It examines the mathematical foundation of game physics and illustrates how it is applied in practice through coding examples. The second half of the book shows you how to use Box2D, a popular open source 2D game physics engine. A companion website provides supplementary material, including source code and videos. This book helps you become a capable 2D game physics programmer through its presentation of both the theory and applications of 2D game physics. After reading the book and experimenting with the code samples, you will understand the basics of 2D game physics and know how to use Box2D to make a 2D physics-based game.

Deep Reinforcement Learning Hands-On

Advanced Information Systems Engineering

Machine Learning and Security

Machine learning avec Scikit-Learn

Concepts, Tools, and Techniques to Build Intelligent Systems

Proceedings of International Conference on Intelligent Computing, Information and Control Systems

Third International Conference, SADASC 2020, Marrakesh, Morocco, June 25–26, 2020, Proceedings

This book outlines the effects that technology-induced change will have on sport within the next five to ten years, and provides food for thought concerning what lies further ahead. Presented as a collection of essays, the authors are leading academics from renowned institutions such as Massachusetts Institute of Technology, Queensland University of Technology, and the University of Cambridge, and practitioners with extensive technological expertise. In their essays, the authors examine the impacts of emerging technologies like artificial intelligence, the Internet of Things, and robotics on sports and assess how they will change sport itself, consumer behavior, and existing business models. The book will help athletes, entrepreneurs, and innovators working in the sports industry to spot trendsetting technologies, gain deeper insights into how they will affect their activities, and identify the most effective responses to stay ahead of the competition both on and off the pitch.

Python est devenu en quelques années un langage majeur dans l'univers des applications centrées sur le traitement des données, et plus particulièrement des gros volumes de données (big data). Cet ouvrage servira de guide à tous ceux qui s'intéressent à l'utilisation de Python pour le travail sur les données et l'automatisation de certaines tâches (data science). Il met l'accent sur la préparation et la mise en forme des données qui sont essentielles dans la qualité du résultat et qui constituent aujourd'hui une part importante du travail du data scientist. L'ensemble des exemples et des exercices présentés dans cet ouvrage sont disponibles sous forme de Notebooks Jupyter. Ils sont accessibles directement sur GitHub dans le répertoire dédié à l'ouvrage ou en téléchargement sur le site Dunod. Cette 2e édition comportent de très nombreuses mises à jour sur les évolutions récentes du langage Python, sur les "packages"

Online Library Deep Learning Avec Tensorflow Mise En Oeuvre Et Cas Concrets

utilisés en data science, et sur d'autres logiciels associés comme TensorFlow et Keras.

Unlock the power of artificial intelligence with top Udemy AI instructor Hadelin de Ponteves. Key Features Learn from friendly, plain English explanations and practical activities Put ideas into action with 5 hands-on projects that show step-by-step how to build intelligent software Use AI to win classic video games and construct a virtual self-driving car Book Description Welcome to the Robot World ... and start building intelligent software now! Through his best-selling video courses, Hadelin de Ponteves has taught hundreds of thousands of people to write AI software. Now, for the first time, his hands-on, energetic approach is available as a book. Starting with the basics before easing you into more complicated formulas and notation, AI Crash Course gives you everything you need to build AI systems with reinforcement learning and deep learning. Five full working projects put the ideas into action, showing step-by-step how to build intelligent software using the best and easiest tools for AI programming, including Python, TensorFlow, Keras, and PyTorch. AI Crash Course teaches everyone to build an AI to work in their applications. Once you've read this book, you're only limited by your imagination. What you will learn Master the basics of AI without any previous experience Build fun projects, including a virtual-self-driving car and a robot warehouse worker Use AI to solve real-world business problems Learn how to code in Python Discover the 5 principles of reinforcement learning Create your own AI toolkit Who this book is for If you want to add AI to your skillset, this book is for you. It doesn't require data science or machine learning knowledge. Just maths basics (high school level).

Providing code examples in python, this book introduces the concepts of machine learning with mathematical explanations and programming fundamentals. --

Intelligent Mobile Projects with TensorFlow

Natural Language Processing with Transformers

Mise en oeuvre et cas concrets

Learning from Data

Des bases du langage au machine learning

Building Recommender Systems with Machine Learning and AI: Help People Discover New Products and Content with Deep Learning, Neural Networks, and Mach

Mise en oeuvre et cas concrets

Machine learning (ML) is changing virtually every aspect of our lives. Today ML algorithms accomplish tasks that until recently only expert humans could perform. As it relates to finance, this is the most exciting time to adopt a disruptive technology that will transform how everyone invests for generations. Readers will learn how to structure Big data in a way that is amenable to ML algorithms; how to conduct research with ML algorithms on that data; how to use supercomputing methods; how to backtest your discoveries while avoiding false positives. The book addresses real-life problems faced by practitioners on a daily basis, and explains scientifically sound solutions using math, supported by code and examples. Readers become active users who can test the proposed solutions in their particular setting. Written by a recognized expert and portfolio manager, this book will equip investment professionals with the groundbreaking tools needed to succeed in modern finance.

This volume constitutes refereed proceedings of the Third International Conference on Smart Applications and Data Analysis, SADASC 2020, held in Marrakesh, Morocco.

Due to the COVID-19 pandemic the conference has been postponed to June 2020. The 24 full papers and 3 short papers presented were thoroughly reviewed and selected from 44 submissions. The papers are organized according to the following topics: ontologies and meta modeling; cyber physical systems and block-chains; recommender systems; machine learning based applications; combinatorial optimization; simulations

and deep learning.

This book gathers papers presented at the second installment of the International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD-2019), which was held on July 08–11, 2019 in Marrakech, Morocco. It offers comprehensive coverage of recent advances in big data, data analytics and related paradigms. The book consists of fifty-two chapters, each of which shares the latest research in the fields of big data and data science, and describes use cases and applications of big data technologies in various domains, such as social networks and health care. All parts of the book discuss open research problems and potential opportunities that have arisen from the rapid advances in big data technologies. In addition, the book surveys the state of the art in data science, and provides practical guidance on big data analytics and data science. Expert perspectives are provided by authoritative researchers and practitioners from around the world, who discuss research developments and emerging trends, present case studies on helpful frameworks and innovative methodologies, and suggest best practices for efficient and effective data analytics. Chiefly intended for researchers, IT professionals and graduate students, the book represents a timely contribution to the growing field of big data, which has been recognized as one of the leading emerging technologies that will have a major impact on various fields of science and various aspects of human society over the next several decades. Therefore, the content in this book is an essential tool to help readers understand current developments, and provides them with an extensive overview of the field of big data analytics as it is practiced today. The chapters cover technical aspects of key areas that generate and use big data, such as management and finance, medicine and health care, networks, the Internet of Things, big data standards, benchmarking of systems, and others. In addition to a diverse range of applications, key algorithmic approaches such as graph partitioning, clustering and finite mixture modeling of high-dimensional data are also covered. The varied collection of topics addressed introduces readers to the richness of the emerging field of big data analytics.

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.

How to Build Applied Machine Learning Solutions from Unlabeled Data
Advanced Intelligent Systems for Sustainable Development (AI2SD ' 2019)
Hands-On Reinforcement Learning with Python
Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow
ICMISC 2020

Apply modern RL methods, with deep Q-networks, value iteration, policy gradients, TRPO, AlphaGo Zero and more

A fun and hands-on introduction to machine learning, reinforcement learning, deep learning, and artificial intelligence with Python

Reinforcement learning is a self-evolving type of machine learning that takes us closer to achieving true artificial intelligence. This easy-to-follow guide explains everything from scratch using rich examples written in Python.

This book is a collection of papers presented at the International Conference on Intelligent Computing, Information and Control Systems (ICICCS 2020). It encompasses various research works that help to develop and advance the next-generation intelligent computing and control systems. The book integrates the computational intelligence and intelligent control systems to provide a powerful methodology for a wide range of data analytics issues in industries and societal applications. The book also presents the new algorithms and methodologies for promoting advances in common intelligent computing and control methodologies including evolutionary computation, artificial life, virtual infrastructures, fuzzy logic, artificial immune systems, neural networks and various neuro-hybrid methodologies. This book is pragmatic for researchers, academicians and students dealing with mathematically intransigent problems. Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and

Online Library Deep Learning Avec Tensorflow Mise En Oeuvre Et Cas Concrets

TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Discover the skill-sets required to implement various approaches to Machine Learning with Python Key FeaturesExplore unsupervised learning with clustering, autoencoders, restricted Boltzmann machines, and moreBuild your own neural network models using modern Python librariesPractical examples show you how to implement different machine learning and deep learning techniquesBook Description

Unsupervised learning is about making use of raw, untagged data and applying learning algorithms to it to help a machine predict its outcome. With this book, you will explore the concept of unsupervised learning to cluster large sets of data and analyze them repeatedly until the desired outcome is found using Python. This book starts with the key differences between supervised, unsupervised, and semi-supervised learning. You will be introduced to the best-used libraries and frameworks from the Python ecosystem and address unsupervised learning in both the machine learning and deep learning domains. You will explore various algorithms, techniques that are used to implement unsupervised learning in real-world use cases. You will learn a variety of unsupervised learning approaches, including randomized optimization, clustering, feature selection and transformation, and information theory. You will get hands-on experience with how neural networks can be employed in unsupervised scenarios. You will also explore the steps involved in building and training a GAN in order to process images. By the end of this book, you will have learned the art of unsupervised learning for different real-world challenges. What you will learnUse cluster algorithms to identify and optimize natural groups of dataExplore advanced non-linear and hierarchical clustering in actionSoft label assignments for fuzzy c-means and Gaussian mixture modelsDetect anomalies through density estimationPerform principal component analysis using neural network modelsCreate unsupervised models using GANsWho this book is for This book is intended for statisticians, data scientists, machine learning developers, and deep learning practitioners who want to build smart applications by implementing key building block unsupervised learning, and master all the new techniques and algorithms offered in machine learning and deep learning using real-world examples. Some prior knowledge of machine learning concepts and statistics is desirable.

Introduction to Game Physics with Box2D

Neural Networks

Programming with TensorFlow

Deep Learning for Computer Vision

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

21st Century Sports

Build 10+ Artificial Intelligence Apps Using TensorFlow Mobile and Lite for IOS, Android, and Raspberry Pi

State-of-the-art algorithms and theory in a novel domain of machine learning, prediction when the output has structure.

Learn how to build recommender systems from one of Amazon's pioneers in the field. Frank Kane spent over nine years at Amazon, where he managed and led the development of many of Amazon's personalized product recommendation technologies. You've seen automated recommendations everywhere - on Netflix's home page, on YouTube, and on Amazon as these machine learning algorithms learn about your unique interests, and show the best products or content for you as an individual. These technologies have become central to the largest, most prestigious tech employers out there, and by understanding how they work, you'll become very valuable to them. This book is adapted from Frank's popular online course published by Sundog Education, so

Online Library Deep Learning Avec Tensorflow Mise En Oeuvre Et Cas Concrets

you can expect lots of visual aids from its slides and a conversational, accessible tone throughout the book. The graphics and scripts from over 300 slides are included, and you'll have access to all of the source code associated with it as well. We'll cover tried and true recommendation algorithms based on neighborhood-based collaborative filtering, and work our way up to more modern techniques including matrix factorization and even deep learning with artificial neural networks. Along the way, you'll learn from Frank's extensive industry experience to understand the real-world challenges you'll encounter when applying these algorithms at large scale and with real-world data. This book is very hands-on; you'll develop your own framework for evaluating and combining many different recommendation algorithms together, and you'll even build your own neural networks using Tensorflow to generate recommendations from real-world movie ratings from real people. We'll cover:

- Building a recommendation engine
- Evaluating recommender systems
- Content-based filtering using item attributes
- Neighborhood-based collaborative filtering with user-based, item-based, and KNN CF
- Model-based methods including matrix factorization and SVD
- Applying deep learning, AI, and artificial neural networks to recommendations
- Session-based recommendations with recursive neural networks
- Scaling to massive data sets with Apache Spark machine learning, Amazon DSSTNE deep learning, and AWS SageMaker with factorization machines
- Real-world challenges and solutions with recommender systems
- Case studies from YouTube and Netflix
- Building hybrid, ensemble recommenders

This comprehensive book takes you all the way from the early days of collaborative filtering, to bleeding-edge applications of deep neural networks and modern machine learning techniques for recommending the best items to every individual user. The coding exercises for this book use the Python programming language. We include an intro to Python if you're new to it, but you'll need some prior programming experience in order to use this book successfully. We also include a short introduction to deep learning, Tensorflow, and Keras if you are new to the field of artificial intelligence, but you'll need to be able to understand new computer algorithms. Dive in, and learn about one of the most interesting and lucrative applications of machine learning and deep learning there is!

This book gathers selected research papers presented at the International Conference on Recent Trends in Machine Learning, IOT, Smart Cities & Applications (ICMISC 2020), held on 29-30 March 2020 at CMR Institute of Technology, Hyderabad, Telangana, India. Discussing current trends in machine learning, Internet of things, and smart cities applications, with a focus on multi-disciplinary research in the area of artificial intelligence and cyber-physical systems, this book is a valuable resource for scientists, research scholars and PG students wanting formulate their research ideas and find the future directions in these areas. Further, it serves as a reference work anyone wishing to understand the latest technologies used by practicing engineers around the globe.

Through a series of recent breakthroughs, deep learning has boosted

Online Library Deep Learning Avec Tensorflow Mise En Oeuvre Et Cas Concrets

the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Volume 4 - Advanced Intelligent Systems for Applied Computing Sciences
Progresses in Artificial Intelligence and Neural Systems
Smart Applications and Data Analysis
Machine Learning, Advances in Computing, Renewable Energy and Communication

Solution for Edge Computing Applications

Python pour le data scientist - 2e éd.

Machine Learning with Python Cookbook

Can machine learning techniques solve our computer security problems and finally put an end to the cat-and-mouse game between attackers and defenders? Or is this hope merely hype? Now you can dive into the science and answer this question for yourself! With this practical guide, you'll explore ways to apply machine learning to security issues such as intrusion detection, malware classification, and network analysis. Machine learning and security specialists Clarence Chio and David Freeman provide a framework for discussing the marriage of these two fields, as well as a toolkit of machine-learning algorithms that you can apply to an array of security problems. This book is ideal for security engineers and data scientists alike. Learn how machine learning has contributed to the success of modern spam filters Quickly detect anomalies, including breaches, fraud, and impending system failure Conduct malware analysis by extracting useful information from computer binaries Uncover attackers within the network by finding patterns inside datasets Examine how attackers exploit consumer-facing websites and app functionality Translate your machine learning algorithms from the lab to production Understand the threat attackers pose to machine learning solutions

Summary Deep Learning with Python introduces the field of deep

learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on

an EC2 GPU instance

This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models

This book constitutes the refereed proceedings of the 32nd International Conference on Advanced Information Systems Engineering, CAiSE 2020, held in Grenoble, France, in June 2020.* The 33 full papers presented in this volume were carefully reviewed and selected from 185 submissions. The book also contains one invited talk in full paper length. The papers were organized in topical sections named: distributed applications; AI and big data in IS; process mining and analysis; requirements and modeling; and information systems engineering. Abstracts on the CAiSE 2020 tutorials can be found in the back matter of the volume. *The conference was held virtually due to the COVID-19 pandemic.

Master reinforcement and deep reinforcement learning using OpenAI Gym and TensorFlow

Implement machine learning and deep learning models using Scikit-Learn, TensorFlow, and more

Practical Deep Learning for Cloud, Mobile, and Edge

Deep learning avec TensorFlow

Image Classification, Object Detection, and Face Recognition in Python

Advances in Financial Machine Learning

Proceedings of International Conference on Recent Trends in Machine Learning, IoT, Smart Cities and Applications

This book provides an overview of the current advances in artificial intelligence and neural

nets. Artificial intelligence (AI) methods have shown great capabilities in modelling, prediction and recognition tasks supporting human-machine interaction. At the same time, the issue of emotion has gained increasing attention due to its relevance in achieving human-like interaction with machines. The real challenge is taking advantage of the emotional characterization of humans' interactions to make computers interfacing with them emotionally and socially credible. The book assesses how and to what extent current sophisticated computational intelligence tools might support the multidisciplinary research on the characterization of appropriate system reactions to human emotions and expressions in interactive scenarios. Discussing the latest recent research trends, innovative approaches and future challenges in AI from interdisciplinary perspectives, it is a valuable resource for researchers and practitioners in academia and industry.

This book provides a thorough overview of the ongoing evolution in the application of artificial intelligence (AI) within healthcare and radiology, enabling readers to gain a deeper insight into the technological background of AI and the impacts of new and emerging technologies on medical imaging. After an introduction on game changers in radiology, such as deep learning technology, the technological evolution of AI in computing science and medical image computing is described, with explanation of basic principles and the types and subtypes of AI. Subsequent sections address the use of imaging biomarkers, the development and validation of AI applications, and various aspects and issues relating to the growing role of big data in radiology. Diverse real-life clinical applications of AI are then outlined for different body parts, demonstrating their ability to add value to daily radiology practices. The concluding section focuses on the impact of AI on radiology and the implications for radiologists, for example with respect to training. Written by radiologists and IT professionals, the book will be of high value for radiologists, medical/clinical physicists, IT specialists, and imaging informatics professionals.

Many industry experts consider unsupervised learning the next frontier in artificial intelligence, one that may hold the key to general artificial intelligence. Since the majority of the world's data is unlabeled, conventional supervised learning cannot be applied. Unsupervised learning, on the other hand, can be applied to unlabeled datasets to discover meaningful patterns buried deep in the data, patterns that may be near impossible for humans to uncover. Author Ankur Patel shows you how to apply unsupervised learning using two simple, production-ready Python frameworks: Scikit-learn and TensorFlow using Keras. With code and hands-on examples, data scientists will identify difficult-to-find patterns in data and gain deeper business insight, detect anomalies, perform automatic feature engineering and selection, and generate synthetic datasets. All you need is programming and some machine learning experience to get started. Compare the strengths and weaknesses of the different machine learning approaches: supervised, unsupervised, and reinforcement learning Set up and manage machine learning projects end-to-end Build an anomaly detection system to catch credit card fraud Clusters users into distinct and homogeneous groups Perform semisupervised learning Develop movie recommender systems using restricted Boltzmann machines Generate synthetic images using generative adversarial networks

Le Deep Learning (apprentissage profond) est un ensemble de techniques avancées du Machine Learning qui reposent principalement sur les réseaux de neurones. Le Deep Learning est au cœur des avancées extraordinaires en intelligence artificielle que l'on a pu observer ces dernières années : reconnaissance de la voix ou des visages, voitures autonomes, traduction automatique, etc. Le Deep Learning est récent et il évolue vite. Ce livre en présente les principales techniques : les réseaux de neurones profonds, capables de modéliser toutes sortes de données, les réseaux de convolution, capables de classifier des images, les segmenter et découvrir les objets ou personnes qui s'y trouvent, les réseaux

récurrents, capables de gérer des séquences telles que des phrases, des séries temporelles, ou encore des vidéos, les Autoencoders qui peuvent découvrir toutes sortes de structures dans des données, de façon non supervisée, et enfin le Reinforcement Learning (apprentissage par renforcement) qui permet de découvrir automatiquement les meilleures actions pour effectuer une tâche (par exemple un robot qui apprend à marcher). Ce livre présente TensorFlow, le framework de Deep Learning open source créé et utilisé par Google. Il est accompagné de Jupyter notebooks (disponibles sur github) qui contiennent tous les exemples de code du livre, afin que le lecteur puisse facilement tester et faire varier les programmes pour mettre en oeuvre ses connaissances.

Machine Learning avec Scikit-Learn

Deep Learning with Python

How Technologies Will Change Sports in the Digital Age

32nd International Conference, CAiSE 2020, Grenoble, France, June 8–12, 2020, Proceedings Opportunities, Applications and Risks

Artificial Intelligence in Medical Imaging

Machine Learning with Python

Neural networks represent a powerful data processing technique that has reached maturity and broad application. When clearly understood and appropriately used, they are a mandatory component in the toolbox of any engineer who wants make the best use of the available data, in order to build models, make predictions, mine data, recognize shapes or signals, etc. Ranging from theoretical foundations to real-life applications, this book is intended to provide engineers and researchers with clear methodologies for taking advantage of neural networks in industrial, financial or banking applications, many instances of which are presented in the book. For the benefit of readers wishing to gain deeper knowledge of the topics, the book features appendices that provide theoretical details for greater insight, and algorithmic details for efficient programming and implementation. The chapters have been written by experts and edited to present a coherent and comprehensive, yet not redundant, practically oriented introduction.

In today's world, deep learning source codes and a plethora of open access geospatial images are readily available and easily accessible. However, most people are missing the educational tools to make use of this resource. Deep Learning for Remote Sensing Images with Open Source Software is the first practical book to introduce deep learning techniques using free open source tools for processing real world remote sensing images. The approaches detailed in this book are generic and can be adapted to suit many different applications for remote sensing image processing, including landcover mapping, forestry, urban studies, disaster mapping, image restoration, etc. Written with practitioners and students in mind, this book helps link together the theory and practical use of existing tools and data to apply deep learning techniques on remote sensing images and data. Specific Features of this Book: The first book that explains how to apply deep learning techniques to public, free available data (Spot-7 and Sentinel-2 images, OpenStreetMap vector data), using open source software (QGIS, Orfeo ToolBox, TensorFlow) Presents approaches suited for real world images and data targeting large scale processing and GIS applications Introduces state of the art deep learning architecture families that can be applied to remote sensing world, mainly for landcover mapping, but also for generic approaches (e.g. image restoration) Suited for deep learning beginners and readers with some GIS knowledge. No coding knowledge is required to learn practical skills. Includes deep learning techniques through many step by step remote sensing data processing exercises.

This book gathers selected papers presented at International Conference on Machine Learning, Advances in Computing, Renewable Energy and Communication (MARC 2020), held in Krishna Engineering College, Ghaziabad, India, during December 17–18, 2020. This book discusses key concepts, challenges, and potential solutions in connection with established and emerging topics in advanced computing, renewable energy, and network communications.

Create Deep Learning and Reinforcement Learning apps for multiple platforms with TensorFlow Key

Features Build TensorFlow-powered AI applications for mobile and embedded devices Learn modern AI topics such as computer vision, NLP, and deep reinforcement learning Get practical insights and exclusive working code not available in the TensorFlow documentation Book Description As a developer, you always need to keep an eye out and be ready for what will be trending soon, while also focusing on what's trending currently. So, what's better than learning about the integration of the best of both worlds, the present and the future? Artificial Intelligence (AI) is widely regarded as the next big thing after mobile, and Google's TensorFlow is the leading open source machine learning framework, the hottest branch of AI. This book covers more than 10 complete iOS, Android, and Raspberry Pi apps powered by TensorFlow and built from scratch, running all kinds of cool TensorFlow models offline on-device: from computer vision, speech and language processing to generative adversarial networks and AlphaZero-like deep reinforcement learning. You'll learn how to use or retrain existing TensorFlow models, build your own models, and develop intelligent mobile apps running those TensorFlow models. You'll learn how to quickly build such apps with step-by-step tutorials and how to avoid many pitfalls in the process with lots of hard-earned troubleshooting tips. What you will learn Classify images with transfer learning Detect objects and their locations Transform pictures with amazing art styles Understand simple speech commands Describe images in natural language Recognize drawing with Convolutional Neural Network and Long Short-Term Memory Predict stock price with Recurrent Neural Network in TensorFlow and Keras Generate and enhance images with generative adversarial networks Build AlphaZero-like mobile game app in TensorFlow and Keras Use TensorFlow Lite and Core ML on mobile Develop TensorFlow apps on Raspberry Pi that can move, see, listen, speak, and learn Who this book is for If you're an iOS/Android developer interested in building and retraining others' TensorFlow models and running them in your mobile apps, or if you're a TensorFlow developer and want to run your new and amazing TensorFlow models on mobile devices, this book is for you. You'll also benefit from this book if you're interested in TensorFlow Lite, Core ML, or TensorFlow on Raspberry Pi.

A Short Course

Proceedings of MARC 2020

Hands-On Unsupervised Learning with Python

The Rust Programming Language (Covers Rust 2018)

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

Methodology and Applications

AI Crash Course

Step-by-step tutorials on deep learning neural networks for computer vision in python with Keras.

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. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. What you will learn Understand the DL context of RL and implement complex DL models Learn the foundation of RL: Markov decision processes Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others 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 Some fluency 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.

Cet ouvrage, conçu pour tous ceux qui souhaitent s'initier au Machine Learning (apprentissage automatique) est la traduction de la première partie du best-seller américain Hands-On Machine Learning with Scikit-Learn & TensorFlow. Il ne requiert que peu de connaissances en mathématiques et présente les fondamentaux du Machine Learning d'une façon très pratique à l'aide de Scikit-Learn qui est l'un des frameworks de ML les plus utilisés actuellement. Des exercices corrigés permettent de s'assurer que l'on a assimilé les concepts et que l'on maîtrise les outils. Des compléments en ligne interactifs sous forme de Jupyter notebooks complètent le livre avec des exemples exécutables. Ce premier titre est complété par un second ouvrage intitulé Deep Learning avec TensorFlow.

Protecting Systems with Data and Algorithms

Deep Learning for Remote Sensing Images with Open Source Software

Transformers for Natural Language Processing

Predicting Structured Data

Hands-On Unsupervised Learning Using Python

Deep Learning avec Keras et TensorFlow

ICICCS 2020