

Deep Learning Basics Github Pages

Starting with an easy introduction to KNIME Analytics Platform, this book will take you through the key features of the platform and cover the advanced and latest deep learning concepts in neural networks. In each chapter, you'll solve real-world case studies based on deep learning networks to spark your creativity for new projects.

This book features selected papers presented at the 2nd International Conference on Advanced Computing Technologies and Applications, held at SVKM's Dwarkadas J. Sanghvi College of Engineering, Mumbai, India, from 28 to 29 February 2020. Covering recent advances in next-generation computing, the book focuses on recent developments in intelligent computing, such as linguistic computing, statistical computing, data computing and ambient applications.

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is

training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural

Download File PDF Deep Learning Basics Github Pages

Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

This two-volume set (CCIS 1045 and CCIS 1046) constitutes the refereed proceedings of the Third International Conference on Advances in Computing and Data Sciences, ICACDS 2019, held in Ghaziabad, India, in April 2019. The 112 full papers were carefully reviewed and selected from 621 submissions. The papers are centered around topics like advanced computing, data sciences, distributed systems organizing principles, development frameworks and environments, software verification and validation, computational complexity and cryptography, machine learning theory, database theory, probabilistic representations.

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory,

Download File PDF Deep Learning Basics Github Pages

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

Take the next step in implementing various common and not-so-common neural networks with Tensorflow 1.x About This Book Skill up and implement tricky neural networks using Google's TensorFlow 1.x An easy-to-follow guide that lets you explore reinforcement learning, GANs, autoencoders, multilayer perceptrons and more. Hands-on recipes to work with Tensorflow on desktop, mobile, and cloud environment Who This Book Is For This book is intended for data analysts, data scientists,

Download File PDF Deep Learning Basics Github Pages

machine learning practitioners and deep learning enthusiasts who want to perform deep learning tasks on a regular basis and are looking for a handy guide they can refer to. People who are slightly familiar with neural networks, and now want to gain expertise in working with different types of neural networks and datasets, will find this book quite useful. What You Will Learn Install TensorFlow and use it for CPU and GPU operations Implement DNNs and apply them to solve different AI-driven problems. Leverage different data sets such as MNIST, CIFAR-10, and Youtube8m with TensorFlow and learn how to access and use them in your code. Use TensorBoard to understand neural network architectures, optimize the learning process, and peek inside the neural network black box. Use different regression techniques for prediction and classification problems Build single and multilayer perceptrons in TensorFlow Implement CNN and RNN in TensorFlow, and use it to solve real-world use cases. Learn how restricted Boltzmann Machines can be used to recommend movies. Understand the implementation of Autoencoders and deep belief networks, and use them for emotion detection. Master the different reinforcement learning methods to implement game playing agents. GANs and their implementation using TensorFlow. In Detail Deep neural networks (DNNs) have achieved a lot of success in the field of computer vision, speech

Download File PDF Deep Learning Basics Github Pages

recognition, and natural language processing. The entire world is filled with excitement about how deep networks are revolutionizing artificial intelligence. This exciting recipe-based guide will take you from the realm of DNN theory to implementing them practically to solve the real-life problems in artificial intelligence domain. In this book, you will learn how to efficiently use TensorFlow, Google's open source framework for deep learning. You will implement different deep learning networks such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Deep Q-learning Networks (DQNs), and Generative Adversarial Networks (GANs) with easy to follow independent recipes. You will learn how to make Keras as backend with TensorFlow. With a problem-solution approach, you will understand how to implement different deep neural architectures to carry out complex tasks at work. You will learn the performance of different DNNs on some popularly used data sets such as MNIST, CIFAR-10, Youtube8m, and more. You will not only learn about the different mobile and embedded platforms supported by TensorFlow but also how to set up cloud platforms for deep learning applications. Get a sneak peek of TPU architecture and how they will affect DNN future. By using crisp, no-nonsense recipes, you will become an expert in implementing deep learning techniques in growing real-world applications and research areas such as

Download File PDF Deep Learning Basics Github Pages

reinforcement learning, GANs, autoencoders and more. Style and approach This book consists of hands-on recipes where you'll deal with real-world problems. You'll execute a series of tasks as you walk through data mining challenges using TensorFlow 1.x. Your one-stop solution for common and not-so-common pain points, this is a book that you must have on the shelf.

This book is focused on the use of deep learning (DL) and artificial intelligence (AI) as tools to advance the fields of malware detection and analysis. The individual chapters of the book deal with a wide variety of state-of-the-art AI and DL techniques, which are applied to a number of challenging malware-related problems. DL and AI based approaches to malware detection and analysis are largely data driven and hence minimal expert domain knowledge of malware is needed. This book fills a gap between the emerging fields of DL/AI and malware analysis. It covers a broad range of modern and practical DL and AI techniques, including frameworks and development tools enabling the audience to innovate with cutting-edge research advancements in a multitude of malware (and closely related) use cases.

This book includes selected papers from the International Conference on Machine Learning and Information Processing (ICMLIP 2019), held at ISB & M School of Technology, Pune, Maharashtra, India,

Download File PDF Deep Learning Basics Github Pages

from December 27 to 28, 2019. It presents the latest developments and technical solutions in the areas of advanced computing and data sciences, covering machine learning, artificial intelligence, human-computer interaction, IoT, deep learning, image processing and pattern recognition, and signal and speech processing.

Machine learning (ML) is the fastest growing field in computer science, and Health Informatics (HI) is amongst the greatest application challenges, providing future benefits in improved medical diagnoses, disease analyses, and pharmaceutical development. However, successful ML for HI needs a concerted effort, fostering integrative research between experts ranging from diverse disciplines from data science to visualization. Tackling complex challenges needs both disciplinary excellence and cross-disciplinary networking without any boundaries. Following the HCI-KDD approach, in combining the best of two worlds, it is aimed to support human intelligence with machine intelligence. This state-of-the-art survey is an output of the international HCI-KDD expert network and features 22 carefully selected and peer-reviewed chapters on hot topics in machine learning for health informatics; they discuss open problems and future challenges in order to stimulate further research and international progress in this field. This guide will help you to explore automated machine learning (AutoML), a rapidly growing subfield of machine learning. You'll learn how you can use AutoML to fully automate the machine learning process even if you're

Download File PDF Deep Learning Basics Github Pages

not an expert, and in turn increase your productivity drastically.

Learn to harness the power of AI for natural language processing, performing tasks such as spell check, text summarization, document classification, and natural language generation. Along the way, you will learn the skills to implement these methods in larger infrastructures to replace existing code or create new algorithms. Applied Natural Language Processing with Python starts with reviewing the necessary machine learning concepts before moving onto discussing various NLP problems. After reading this book, you will have the skills to apply these concepts in your own professional environment. What You Will Learn Utilize various machine learning and natural language processing libraries such as TensorFlow, Keras, NLTK, and Gensim Manipulate and preprocess raw text data in formats such as .txt and .pdf Strengthen your skills in data science by learning both the theory and the application of various algorithms Who This Book Is For You should be at least a beginner in ML to get the most out of this text, but you needn't feel that you need be an expert to understand the content.

DEEP LEARNING FOR THE EARTH SCIENCES

Explore this insightful treatment of deep learning in the field of earth sciences, from four leading voices Deep learning is a fundamental technique in modern Artificial Intelligence and is being applied to disciplines across the scientific spectrum; earth science is no exception. Yet, the link between deep learning and Earth sciences has only recently entered academic curricula and thus has

Download File PDF Deep Learning Basics Github Pages

not yet proliferated. Deep Learning for the Earth Sciences delivers a unique perspective and treatment of the concepts, skills, and practices necessary to quickly become familiar with the application of deep learning techniques to the Earth sciences. The book prepares readers to be ready to use the technologies and principles described in their own research. The distinguished editors have also included resources that explain and provide new ideas and recommendations for new research especially useful to those involved in advanced research education or those seeking PhD thesis orientations. Readers will also benefit from the inclusion of: An introduction to deep learning for classification purposes, including advances in image segmentation and encoding priors, anomaly detection and target detection, and domain adaptation An exploration of learning representations and unsupervised deep learning, including deep learning image fusion, image retrieval, and matching and co-registration Practical discussions of regression, fitting, parameter retrieval, forecasting and interpolation An examination of physics-aware deep learning models, including emulation of complex codes and model parametrizations Perfect for PhD students and researchers in the fields of geosciences, image processing, remote sensing, electrical engineering and computer science, and machine learning, Deep Learning for the Earth Sciences will also earn a place in the libraries of machine learning and pattern recognition researchers, engineers, and scientists.

Examine the latest technological advancements in

Download File PDF Deep Learning Basics Github Pages

building a scalable machine-learning model with big data using R. This second edition shows you how to work with a machine-learning algorithm and use it to build a ML model from raw data. You will see how to use R programming with TensorFlow, thus avoiding the effort of learning Python if you are only comfortable with R. As in the first edition, the authors have kept the fine balance of theory and application of machine learning through various real-world use-cases which gives you a comprehensive collection of topics in machine learning. New chapters in this edition cover time series models and deep learning. What You'll Learn Understand machine learning algorithms using R Master the process of building machine-learning models Cover the theoretical foundations of machine-learning algorithms See industry focused real-world use cases Tackle time series modeling in R Apply deep learning using Keras and TensorFlow in R Who This Book is For Data scientists, data science professionals, and researchers in academia who want to understand the nuances of machine-learning approaches/algorithms in practice using R.

Statistics, Data Mining, and Machine Learning in Astronomy is the essential introduction to the statistical methods needed to analyze complex data sets from astronomical surveys such as the Panoramic Survey Telescope and Rapid Response System, the Dark Energy Survey, and the Large Synoptic Survey Telescope. Now fully updated, it presents a wealth of practical analysis problems, evaluates the techniques for solving them, and explains how to use various

Download File PDF Deep Learning Basics Github Pages

approaches for different types and sizes of data sets. Python code and sample data sets are provided for all applications described in the book. The supporting data sets have been carefully selected from contemporary astronomical surveys and are easy to download and use. The accompanying Python code is publicly available, well documented, and follows uniform coding standards. Together, the data sets and code enable readers to reproduce all the figures and examples, engage with the different methods, and adapt them to their own fields of interest. An accessible textbook for students and an indispensable reference for researchers, this updated edition features new sections on deep learning methods, hierarchical Bayes modeling, and approximate Bayesian computation. The chapters have been revised throughout and the astroML code has been brought completely up to date. Fully revised and expanded

Describes the most useful statistical and data-mining methods for extracting knowledge from huge and complex astronomical data sets Features real-world data sets from astronomical surveys Uses a freely available Python codebase throughout Ideal for graduate students, advanced undergraduates, and working astronomers

Learn basic Python programming to create functional and effective visualizations from earth observation satellite data sets Thousands of satellite datasets are freely available online, but scientists need the right tools to efficiently analyze data and share results. Python has easy-to-learn syntax and thousands of libraries to perform common Earth science programming tasks.

Earth Observation Using Python: A Practical

Download File PDF Deep Learning Basics Github Pages

Programming Guide presents an example-driven collection of basic methods, applications, and visualizations to process satellite data sets for Earth science research. Gain Python fluency using real data and case studies Read and write common scientific data formats, like netCDF, HDF, and GRIB2 Create 3-dimensional maps of dust, fire, vegetation indices and more Learn to adjust satellite imagery resolution, apply quality control, and handle big files Develop useful workflows and learn to share code using version control Acquire skills using online interactive code available for all examples in the book The American Geophysical Union promotes discovery in Earth and space science for the benefit of humanity. Its publications disseminate scientific knowledge and provide resources for researchers, students, and professionals. Find out more about this book from this Q&A with the Author 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 TensorFlow Lite Develop AI for a range of

Download File PDF Deep Learning Basics Github Pages

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 Use Java and Deeplearning4j to build robust, scalable, and highly accurate AI models from scratch Key Features Install and configure Deeplearning4j to implement deep learning models from scratch Explore recipes for developing, training, and fine-tuning your neural network models in Java Model neural networks using datasets containing images, text, and time-series data Book Description Java is one of the most widely used programming languages in the world. With this book, you will see how to perform deep learning using Deeplearning4j (DL4J) – the most popular Java library for training neural networks efficiently. This book starts by showing you how to install and configure Java and DL4J on your system. You will then gain insights into deep learning basics and use your knowledge to create a deep neural network for binary classification from scratch. As you progress, you will discover how to build a convolutional neural network (CNN) in DL4J, and understand how to construct numeric vectors from text. This deep learning book will also guide you through performing anomaly detection on unsupervised data and help you set up neural networks in distributed systems effectively. In addition to this, you will learn how to import

Download File PDF Deep Learning Basics Github Pages

models from Keras and change the configuration in a pre-trained DL4J model. Finally, you will explore benchmarking in DL4J and optimize neural networks for optimal results. By the end of this book, you will have a clear understanding of how you can use DL4J to build robust deep learning applications in Java. What you will learn

- Perform data normalization and wrangling using DL4J
- Build deep neural networks using DL4J
- Implement CNNs to solve image classification problems
- Train autoencoders to solve anomaly detection problems using DL4J
- Perform benchmarking and optimization to improve your model's performance
- Implement reinforcement learning for real-world use cases using RL4J
- Leverage the capabilities of DL4J in distributed systems

Who this book is for If you are a data scientist, machine learning developer, or a deep learning enthusiast who wants to implement deep learning models in Java, this book is for you. Basic understanding of Java programming as well as some experience with machine learning and neural networks is required to get the most out of this book. Linear algebra is a pillar of machine learning. You cannot develop a deep understanding and application of machine learning without it. In this laser-focused Ebook, you will finally cut through the equations, Greek letters, and confusion, and discover the topics in linear algebra that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover what linear algebra is, the importance of linear algebra to machine learning, vector, and matrix operations, matrix factorization, principal component analysis, and much more.

Download File PDF Deep Learning Basics Github Pages

This book is a part of the Blue Book series “Research on the Development of Electronic Information Engineering Technology in China,” which explores the cutting edge of deep learning studies. A subfield of machine learning, deep learning differs from conventional machine learning methods in its ability to learn multiple levels of representation and abstraction by using several layers of nonlinear modules for feature extraction and transformation. The extensive use and huge success of deep learning in speech, CV, and NLP have led to significant advances toward the full materialization of AI. Focusing on the development of deep learning technologies, this book also discusses global trends, the status of deep learning development in China and the future of deep learning.

Learn the basics of analytics on big data using Java, machine learning and other big data tools About This Book Acquire real-world set of tools for building enterprise level data science applications Surpasses the barrier of other languages in data science and learn create useful object-oriented codes Extensive use of Java compliant big data tools like apache spark, Hadoop, etc. Who This Book Is For This book is for Java developers who are looking to perform data analysis in production environment. Those who wish to implement data analysis in their Big data applications will find this book helpful. What You Will Learn Start from simple analytic tasks on big data Get into more complex tasks with predictive analytics on big data using machine learning Learn real time analytic tasks Understand the concepts with examples and case studies Prepare and

Download File PDF Deep Learning Basics Github Pages

refine data for analysis Create charts in order to understand the data See various real-world datasets In Detail This book covers case studies such as sentiment analysis on a tweet dataset, recommendations on a movielens dataset, customer segmentation on an ecommerce dataset, and graph analysis on actual flights dataset. This book is an end-to-end guide to implement analytics on big data with Java. Java is the de facto language for major big data environments, including Hadoop. This book will teach you how to perform analytics on big data with production-friendly Java. This book basically divided into two sections. The first part is an introduction that will help the readers get acquainted with big data environments, whereas the second part will contain a hardcore discussion on all the concepts in analytics on big data. It will take you from data analysis and data visualization to the core concepts and advantages of machine learning, real-life usage of regression and classification using Naive Bayes, a deep discussion on the concepts of clustering, and a review of simple neural networks on big data using deepLearning4j or plain Java Spark code. This book is a must-have book for Java developers who want to start learning big data analytics and want to use it in the real world. Style and approach The approach of book is to deliver practical learning modules in manageable content. Each chapter is a self-contained unit of a concept in big data analytics. Book will step by step builds the competency in the area of big data analytics. Examples using real world case studies to give ideas of real applications and how to use the techniques mentioned. The examples and case

Download File PDF Deep Learning Basics Github Pages

studies will be shown using both theory and code.

If machine learning transforms the nature of knowledge, does it also transform the practice of critical thought?

Machine learning—programming computers to learn from data—has spread across scientific disciplines, media, entertainment, and government. Medical research, autonomous vehicles, credit transaction processing, computer gaming, recommendation systems, finance, surveillance, and robotics use machine learning.

Machine learning devices (sometimes understood as scientific models, sometimes as operational algorithms) anchor the field of data science. They have also become mundane mechanisms deeply embedded in a variety of systems and gadgets. In contexts from the everyday to the esoteric, machine learning is said to transform the nature of knowledge. In this book, Adrian Mackenzie investigates whether machine learning also transforms the practice of critical thinking. Mackenzie focuses on machine learners—either humans and machines or human-machine relations—situated among settings, data, and devices. The settings range from fMRI to Facebook; the data anything from cat images to DNA sequences; the devices include neural networks, support vector machines, and decision trees. He examines specific learning algorithms—writing code and writing about code—and develops an archaeology of operations that, following Foucault, views machine learning as a form of knowledge production and a strategy of power. Exploring layers of abstraction, data infrastructures, coding practices, diagrams, mathematical formalisms, and the social organization of machine learning, Mackenzie

Download File PDF Deep Learning Basics Github Pages

traces the mostly invisible architecture of one of the central zones of contemporary technological cultures. Mackenzie's account of machine learning locates places in which a sense of agency can take root. His archaeology of the operational formation of machine learning does not unearth the footprint of a strategic monolith but reveals the local tributaries of force that feed into the generalization and plurality of the field. This book constitutes the refereed proceedings of the 6th Latin American High Performance Computing Conference, CARLA 2019, held in Turrialba, Costa Rica, in September 2019. The 32 revised full papers presented were carefully reviewed and selected out of 62 submissions. The papers included in this book are organized according to the conference tracks - regular track on high performance computing: applications; algorithms and models; architectures and infrastructures; and special track on bioinspired processing (BIP): neural and evolutionary approaches; image and signal processing; biodiversity informatics and computational biology.

Get savvy with R language and actualize projects aimed at analysis, visualization and machine learning About This Book Proficiently analyze data and apply machine learning techniques Generate visualizations, develop interactive visualizations and applications to understand various data exploratory functions in R Construct a predictive model by using a variety of machine learning packages Who This Book Is For This Learning Path is ideal for those who have been exposed to R, but have not used it extensively yet. It covers the basics of using

Download File PDF Deep Learning Basics Github Pages

R and is written for new and intermediate R users interested in learning. This Learning Path also provides in-depth insights into professional techniques for analysis, visualization, and machine learning with R – it will help you increase your R expertise, regardless of your level of experience. What You Will Learn Get data into your R environment and prepare it for analysis Perform exploratory data analyses and generate meaningful visualizations of the data Generate various plots in R using the basic R plotting techniques Create presentations and learn the basics of creating apps in R for your audience Create and inspect the transaction dataset, performing association analysis with the Apriori algorithm Visualize associations in various graph formats and find frequent itemset using the ECLAT algorithm Build, tune, and evaluate predictive models with different machine learning packages Incorporate R and Hadoop to solve machine learning problems on big data In Detail The R language is a powerful, open source, functional programming language. At its core, R is a statistical programming language that provides impressive tools to analyze data and create high-level graphics. This Learning Path is chock-full of recipes. Literally! It aims to excite you with awesome projects focused on analysis, visualization, and machine learning. We'll start off with data analysis – this will show you ways to use R to generate professional analysis reports. We'll then move on to visualizing our data – this provides you with all the guidance needed to get comfortable with data visualization with R. Finally, we'll move into the world of machine learning – this introduces you to data

Download File PDF Deep Learning Basics Github Pages

classification, regression, clustering, association rule mining, and dimension reduction. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: R Data Analysis Cookbook by Viswa Viswanathan and Shanthi Viswanathan R Data Visualization Cookbook by Atmajitsinh Gohil Machine Learning with R Cookbook by Yu-Wei, Chiu (David Chiu) Style and approach This course creates a smooth learning path that will teach you how to analyze data and create stunning visualizations. The step-by-step instructions provided for each recipe in this comprehensive Learning Path will show you how to create machine learning projects with R.

We're living in a digital world. Most of our global economy is digital and the sheer volume of data is stupendous. It's 2020 and we're living in the future. Data Scientist is one of the hottest job on the market right now. Demand for data science is huge and will only grow, and it seems like it will grow much faster than the actual number of data scientists. So if you want to make a career change and become a data scientist, now is the time. This book will guide you through the process. From my experience of working with multiple companies as a project manager, a data science consultant or a CTO, I was able to see the process of hiring data scientists and building data science teams. I know what's important to land your first job as a data scientist, what skills you should acquire, what you should show during a job interview.

This book presents a wealth of deep-learning algorithms

Download File PDF Deep Learning Basics Github Pages

and demonstrates their design process. It also highlights the need for a prudent alignment with the essential characteristics of the nature of learning encountered in the practical problems being tackled. Intended for readers interested in acquiring practical knowledge of analysis, design, and deployment of deep learning solutions to real-world problems, it covers a wide range of the paradigm's algorithms and their applications in diverse areas including imaging, seismic tomography, smart grids, surveillance and security, and health care, among others. Featuring systematic and comprehensive discussions on the development processes, their evaluation, and relevance, the book offers insights into fundamental design strategies for algorithms of deep learning.

Take your Python machine learning ideas and create serverless web applications accessible by anyone with an Internet connection. Some of the most popular serverless cloud providers are covered in this book—Amazon, Microsoft, Google, and PythonAnywhere. You will work through a series of common Python data science problems in an increasing order of complexity. The practical projects presented in this book are simple, clear, and can be used as templates to jump-start many other types of projects. You will learn to create a web application around numerical or categorical predictions, understand the analysis of text, create powerful and interactive presentations, serve restricted access to data, and leverage web plugins to accept credit card payments and donations. You will get your projects into the hands of the world in no time. Each chapter follows three steps:

Download File PDF Deep Learning Basics Github Pages

modeling the right way, designing and developing a local web application, and deploying onto a popular and reliable serverless cloud provider. You can easily jump to or skip particular topics in the book. You also will have access to Jupyter notebooks and code repositories for complete versions of the code covered in the book. What You'll Learn

- Extend your machine learning models using simple techniques to create compelling and interactive web dashboards
- Leverage the Flask web framework for rapid prototyping of your Python models and ideas
- Create dynamic content powered by regression coefficients, logistic regressions, gradient boosting machines, Bayesian classifications, and more
- Harness the power of TensorFlow by exporting saved models into web applications
- Create rich web dashboards to handle complex real-time user input with JavaScript and Ajax to yield interactive and tailored content
- Create dashboards with paywalls to offer subscription-based access
- Access API data such as Google Maps, OpenWeather, etc.
- Apply different approaches to make sense of text data and return customized intelligence
- Build an intuitive and useful recommendation site to add value to users and entice them to keep coming back
- Utilize the freemium offerings of Google Analytics and analyze the results
- Take your ideas all the way to your customer's plate using the top serverless cloud providers

Who This Book Is For

Those with some programming experience with Python, code editing, and access to an interpreter in working order. The book is geared toward entrepreneurs who want to get their ideas onto the web without breaking the bank, small companies without an IT staff,

Download File PDF Deep Learning Basics Github Pages

students wanting exposure and training, and for all data science professionals ready to take things to the next level.

Artificial intelligence is increasingly finding its way into industrial and manufacturing contexts. The prevalence of AI in industry from stock market trading to manufacturing makes it easy to forget how complex artificial intelligence has become. Engineering provides various current and prospective applications of these new and complex artificial intelligence technologies. Applications of Artificial Intelligence in Electrical Engineering is a critical research book that examines the advancing developments in artificial intelligence with a focus on theory and research and their implications. Highlighting a wide range of topics such as evolutionary computing, image processing, and swarm intelligence, this book is essential for engineers, manufacturers, technology developers, IT specialists, managers, academicians, researchers, computer scientists, and students.

This book is a survey and analysis of how deep learning can be used to generate musical content. The authors offer a comprehensive presentation of the foundations of deep learning techniques for music generation. They also develop a conceptual framework used to classify and analyze various types of architecture, encoding models, generation strategies, and ways to control the generation. The five dimensions of this framework are: objective (the kind of musical content to be generated, e.g., melody, accompaniment); representation (the musical elements to be considered and how to encode them, e.g., chord, silence, piano roll, one-hot encoding);

Download File PDF Deep Learning Basics Github Pages

architecture (the structure organizing neurons, their connexions, and the flow of their activations, e.g., feedforward, recurrent, variational autoencoder); challenge (the desired properties and issues, e.g., variability, incrementality, adaptability); and strategy (the way to model and control the process of generation, e.g., single-step feedforward, iterative feedforward, decoder feedforward, sampling). To illustrate the possible design decisions and to allow comparison and correlation analysis they analyze and classify more than 40 systems, and they discuss important open challenges such as interactivity, originality, and structure. The authors have extensive knowledge and experience in all related research, technical, performance, and business aspects. The book is suitable for students, practitioners, and researchers in the artificial intelligence, machine learning, and music creation domains. The reader does not require any prior knowledge about artificial neural networks, deep learning, or computer music. The text is fully supported with a comprehensive table of acronyms, bibliography, glossary, and index, and supplementary material is available from the authors' website.

Summary Companies small and large are initiating AI projects, investing vast sums of money on software, developers, and data scientists. Too often, these AI projects focus on technology at the expense of actionable or tangible business results, resulting in scattershot results and wasted investment. Succeeding with AI sets out a blueprint for AI projects to ensure they are predictable, successful, and profitable. It's filled with practical techniques for running data science programs

Download File PDF Deep Learning Basics Github Pages

that ensure they're cost effective and focused on the right business goals. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Succeeding with AI requires talent, tools, and money. So why do many well-funded, state-of-the-art projects fail to deliver meaningful business value? Because talent, tools, and money aren't enough: You also need to know how to ask the right questions. In this unique book, AI consultant Veljko Kronic reveals a tested process to start AI projects right, so you'll get the results you want. About the book Succeeding with AI sets out a framework for planning and running cost-effective, reliable AI projects that produce real business results. This practical guide reveals secrets forged during the author's experience with dozens of startups, established businesses, and Fortune 500 giants that will help you establish meaningful, achievable goals. In it you'll master a repeatable process to maximize the return on data-scientist hours and learn to implement effectiveness metrics for keeping projects on track and resistant to calcification. What's inside Where to invest for maximum payoff How AI projects are different from other software projects Catching early warnings in time to correct course Exercises and examples based on real-world business dilemmas About the reader For project and business leadership, result-focused data scientists, and engineering teams. No AI knowledge required. About the author Veljko Kronic is a data science consultant, has a computer science PhD, and is a certified Six Sigma Master Black Belt. Table of Contents: 1. Introduction 2.

Download File PDF Deep Learning Basics Github Pages

How to use AI in your business 3. Choosing your first AI project 4. Linking business and technology 5. What is an ML pipeline, and how does it affect an AI project? 6. Analyzing an ML pipeline 7. Guiding an AI project to success 8. AI trends that may affect you

A coherent introduction to core concepts and deep learning techniques that are critical to academic research and real-world applications.

The two-volume set LNCS 11944-11945 constitutes the proceedings of the 19th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2019, held in Melbourne, Australia, in December 2019. The 73 full and 29 short papers presented were carefully reviewed and selected from 251 submissions.

The papers are organized in topical sections on: Parallel and Distributed Architectures, Software Systems and Programming Models, Distributed and Parallel and Network-based Computing, Big Data and its Applications, Distributed and Parallel Algorithms, Applications of Distributed and Parallel Computing, Service Dependability and Security, IoT and CPS Computing, Performance Modelling and Evaluation.

KEY FEATURES ? Understand applications like reinforcement learning, automatic driving and image generation. ? Understand neural networks accompanied with figures and charts. ? Learn about determining coefficients and initial values of weights. **DESCRIPTION**

Deep learning helps you solve issues related to data problems as it has a vast array of mathematical algorithms and has capacity to detect patterns. This book starts with a quick view of deep learning in Python which

Download File PDF Deep Learning Basics Github Pages

would include definition, features and applications. You would be learning about perceptron, neural networks, Backpropagation. This book would also give you a clear insight of how to use Numpy and Matplotlib in deep learning models. By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning. WHAT YOU WILL LEARN ? To develop deep learning applications, use Python with few outside inputs. ? Study several ideas of profound learning and neural networks ? Learn how to determine coefficients of learning and weight values ? Explore applications such as automation, image generation and reinforcement learning ? Implement trends like batch Normalisation, dropout, and Adam WHO THIS BOOK IS FOR Deep Learning from the Basics is for data scientists, data analysts and developers who wish to build efficient solutions by applying deep learning techniques. Individuals who would want a better grasp of technology and an overview. You should have a workable Python knowledge is a required. NumPy knowledge and pandas will be an advantage, but that's completely optional. TABLE OF CONTENTS 1. Python Introduction 2. Perceptron in Depth 3. Neural Networks 4. Training Neural Network 5. Backpropagation 6. Neural Network Training Techniques 7. CNN 8. Deep Learning 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

Download File PDF Deep Learning Basics Github Pages

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. Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively

Download File PDF Deep Learning Basics Github Pages

Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description TensorFlow.js is a framework that enables you to create performant machine learning (ML) applications that run smoothly in a web browser. With this book, you will learn how to use TensorFlow.js to implement various ML models through an example-based approach. Starting with the basics, you'll understand how ML models can be built on the web. Moving on, you will get to grips with the TensorFlow.js ecosystem to develop applications more efficiently. The book will then guide you through implementing ML techniques and algorithms such as regression, clustering, fast Fourier transform (FFT), and dimensionality reduction. You will later cover the Bellman equation to solve Markov decision process (MDP) problems and understand how it is related to reinforcement learning. Finally, you will explore techniques for deploying ML-based web applications and training models with TensorFlow Core. Throughout this ML book, you'll discover useful tips and tricks that will build on your knowledge. By the end of this book, you will be equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset Deploy tf.js-converter to convert Keras models and load them into TensorFlow.js Apply the Bellman equation to solve MDP problems Use the k-means algorithm in TensorFlow.js to visualize prediction results Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps Implement tf.js backend frameworks to tune and accelerate app performance Who this book is for This book is for web developers who

Download File PDF Deep Learning Basics Github Pages

want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started. Discover ways to implement various deep learning algorithms by leveraging Python and other technologies

Key Features

- Learn deep learning models through several activities
- Begin with simple machine learning problems, and finish by building a complex system of your own
- Teach your machines to see by mastering the technologies required for image recognition

Book Description Deep learning is rapidly becoming the most preferred way of solving data problems. This is thanks, in part, to its huge variety of mathematical algorithms and their ability to find patterns that are otherwise invisible to us. Deep Learning from the Basics begins with a fast-paced introduction to deep learning with Python, its definition, characteristics, and applications. You'll learn how to use the Python interpreter and the script files in your applications, and utilize NumPy and Matplotlib in your deep learning models. As you progress through the book, you'll discover backpropagation—an efficient way to calculate the gradients of weight parameters—and study multilayer perceptrons and their limitations, before, finally, implementing a three-layer neural network and calculating multidimensional arrays. By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning. What you will learn

- Use Python with minimum external sources to implement deep learning programs
- Study the various deep learning and neural network theories
- Learn how to determine learning coefficients and the initial values of weights
- Implement trends such as Batch Normalization, Dropout, and Adam
- Explore applications like

Download File PDF Deep Learning Basics Github Pages

automatic driving, image generation, and reinforcement learning Who this book is for Deep Learning from the Basics is designed for data scientists, data analysts, and developers who want to use deep learning techniques to develop efficient solutions. This book is ideal for those who want a deeper understanding as well as an overview of the technologies. Some working knowledge of Python is a must. Knowledge of NumPy and pandas will be beneficial, but not essential. Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies,

Download File PDF Deep Learning Basics Github Pages

music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn

- Execute end-to-end machine learning projects and systems
- Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks
- Review case studies depicting applications of machine learning and deep learning on diverse domains and industries
- Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning.

Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

[Copyright: c36a1ece2b8f9797ff857fc0fafcdf2f](https://c36a1ece2b8f9797ff857fc0fafcdf2f)