

Chatbot Fundamentals An Interactive

Explore the adoption of chatbots in business by focusing on the design, deployment, and continuous improvement of chatbots in a business, with a single use-case from the banking and insurance sector. This book starts by identifying the business processes in the banking and insurance industry. This involves data collection from sources such as conversations from customer service centers, online chats, emails, and other NLP sources. You'll then design the solution architecture of the chatbot. Once the architecture is framed, the author goes on to explain natural language understanding (NLU), natural language processing (NLP), and natural language generation (NLG) with examples. In the next sections, you'll design and implement the backend framework of a typical chatbot from scratch. You will also explore some popular open-source chatbot frameworks such as Dialogflow and LUIS. The authors then explain how you can integrate various third-party services and enterprise databases with the custom chatbot framework. In the final section, you'll discuss how to deploy the custom chatbot framework on the AWS cloud. By the end of Building an Enterprise Chatbot, you will be able to design and develop an enterprise-ready conversational chatbot using an open source development platform to serve the end user. What You Will Learn Identify business processes where chatbots could be used Focus on building a chatbot for one industry and one use-case rather than building a ubiquitous and generic chatbot Design the solution architecture for a chatbot Integrate chatbots with internal data sources using APIs Discover the differences between natural language understanding (NLU), natural language processing (NLP), and natural language generation (NLG) Work with deployment and continuous improvement through representational learning Who This Book Is For Data scientists and enterprise architects who are currently looking to deploy chatbot solutions to their business.

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

Based on the bestselling Marketing by Baines, Fill, Rosengren, and Antonetti, Fundamentals of Marketing is the most relevant, concise guide to marketing, combining the most essential theories with a global range of practitioner insights.

What's the secret to sales success? If you're like most business leaders, you'd say it's fundamentally about relationships-and you'd be wrong. The best salespeople don't just build relationships with customers. They challenge them. The need to understand what top-performing reps are doing that their average performing colleagues are not drove Matthew Dixon, Brent Adamson, and their colleagues at Corporate Executive Board to investigate the skills, behaviors, knowledge, and attitudes that matter most for high performance. And what they discovered may be the biggest shock to conventional sales wisdom in decades. Based on an exhaustive study of thousands of sales reps across multiple industries and geographies, The Challenger Sale argues that classic relationship building is a losing approach, especially when it comes to selling complex, large-scale business-to-business solutions. The authors' study found that every sales rep in the world falls into one of five distinct profiles, and while all of these types of reps can deliver average sales performance, only one-the Challenger- delivers consistently high performance. Instead of bludgeoning customers with endless facts and features about their company and products, Challengers approach customers with unique insights about how they can save or make money. They tailor their sales message to the customer's specific needs and objectives. Rather than acquiescing to the customer's every demand or objection, they are assertive, pushing back when necessary and taking control of the sale. The things that make Challengers unique are replicable and teachable to the average sales rep. Once you understand how to identify the Challengers in your organization, you can model their approach and embed it throughout your sales force. The authors explain how almost any average-performing rep, once equipped with the right tools, can successfully reframe customers' expectations and deliver a distinctive purchase experience that drives higher levels of customer loyalty and, ultimately, greater growth.

This is the only comprehensive, authoritative guide to building Conversational User Interfaces (CUI, a.k.a. bots, chatbots, or chatterbots) with the Microsoft Bot Framework. Reflecting the next radical revolution in human-computer interaction, it will help you leverage advanced artificial intelligence (AI) and natural language processing to empower new and existing applications with stunningly intuitive conversational interfaces. Long-time Microsoft MVP Joe Mayo begins with high-level explanations of what Microsoft Bot Framework is, what you can do with it, and why it matters so much. Next, he presents the foundational knowledge you need to start creating real bots and CUIs. Step by step, you'll learn how to build message dialogs, manage conversations, interact with framework APIs, and incorporate powerful natural language processing with Microsoft's advanced Language Understanding Intelligent Service (LUIS). Mayo also offers detailed guidance on deploying your customized bots to key platforms such as Slack, Skype, and Facebook Messenger. Throughout, Mayo's practical examples combine code with clear explanations of when and why you would perform each task. From start to finish, Programming the Microsoft Bot Framework is relentlessly practical, helping you translate the advanced "magic" of intelligent bots into real solutions right now.

Take a deep dive into the concepts of machine learning as they apply to contemporary business and management. You will learn how machine learning techniques are used to solve fundamental and complex problems in society and industry. Machine Learning for Decision Makers serves as an excellent resource for establishing the relationship of machine learning with IoT, big data, and cognitive and cloud computing to give you an overview of how these modern areas of computing relate to each other. This book introduces a collection of the most important concepts of machine learning and sets them in context with other vital technologies that decision makers need to know about. These concepts span the process from envisioning the problem to applying machine-learning techniques to your particular situation. This discussion also provides an insight to help deploy the results to improve decision-making. The book uses case studies and jargon busting to help you grasp the theory of machine learning quickly. You'll soon gain the big picture of machine learning and how it fits with other cutting-edge IT services. This knowledge will give you confidence in your decisions for the future of your business. What You Will Learn Discover the machine learning, big data, and cloud and cognitive computing technology stack Gain insights into machine learning concepts and practices Understand business and enterprise decision-making using machine learning Absorb machine-learning best practices Who This Book Is For Managers tasked with making key decisions who want to learn how and when machine learning and related technologies can help them. Create AI applications in Python and lay the foundations for your career in data science Key Features Practical examples that explain key machine learning algorithms Explore neural networks in detail with interesting examples Master core AI concepts with engaging activities Book Description Machine learning and neural networks are pillars on which you can build intelligent applications. Artificial Intelligence and Machine Learning Fundamentals begins by introducing you to Python and discussing AI search algorithms. You will cover in-depth

mathematical topics, such as regression and classification, illustrated by Python examples. As you make your way through the book, you will progress to advanced AI techniques and concepts, and work on real-life datasets to form decision trees and clusters. You will be introduced to neural networks, a powerful tool based on Moore's law. By the end of this book, you will be confident when it comes to building your own AI applications with your newly acquired skills! What you will learn Understand the importance, principles, and fields of AI Implement basic artificial intelligence concepts with Python Apply regression and classification concepts to real-world problems Perform predictive analysis using decision trees and random forests Carry out clustering using the k-means and mean shift algorithms Understand the fundamentals of deep learning via practical examples Who this book is for Artificial Intelligence and Machine Learning Fundamentals is for software developers and data scientists who want to enrich their projects with machine learning. You do not need any prior experience in AI. However, it's recommended that you have knowledge of high school-level mathematics and at least one programming language (preferably Python). Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

This book introduces social manufacturing, the next generation manufacturing paradigm that covers product life cycle activities that deal with Internet-based organizational and interactive mechanisms under the context of socio-technical systems in the fields of industrial and production engineering. Like its subject, the book's approach is multi-disciplinary, including manufacturing systems, operations management, computational social sciences and information systems applications. It reports on the latest research findings regarding the social manufacturing paradigm, the architecture, configuration and execution of social manufacturing systems and more. Further, it describes the individual technologies enabled by social manufacturing for each topic, supported by case studies. The technologies discussed include manufacturing resource minimalization and their socialized reorganizations, blockchain models in cybersecurity, computing and decision-making, social business relationships and organizational networks, open product design, social sensors and extended cyber-physical systems, and social factory and inter-connections. This book helps engineers and managers in industry to practice social manufacturing, as well as offering a systematic reference resource for researchers in manufacturing. Students also benefit from the detailed discussions of the latest research and technologies that will have been put into practice by the time they graduate.

Learn some of Python's most popular libraries like BeautifulSoup, Matplotlib, Numpy, Requests, and Tkinter by scraping a website for aluminum alloy information, graphing satellite orbits, creating a GUI to convert different units, plotting airfoil coordinates, publishing programs to PDF, and more!

Using a web API to provide services to application developers is one of the more satisfying endeavors that software engineers undertake. But building a popular API with a thriving developer ecosystem is also one of the most challenging. With this practical guide, developers, architects, and tech leads will learn how to navigate complex decisions for designing, scaling, marketing, and evolving interoperable APIs. Authors Brenda Jin, Saurabh Sahni, and Amir Shevat explain API design theory and provide hands-on exercises for building your web API and managing its operation in production. You'll also learn how to build and maintain a following of app developers. This book includes expert advice, worksheets, checklists, and case studies from companies including Slack, Stripe, Facebook, Microsoft, Cloudinary, Oracle, and GitHub. Get an overview of request-response and event-driven API design paradigms Learn best practices for designing an API that meets the needs of your users Use a template to create an API design process Scale your web API to support a growing number of API calls and use cases Regularly adapt the API to reflect changes to your product or business Provide developer resources that include API documentation, samples, and tools

Packed with intriguing real-world projects as well as theory, Generative AI with Python and TensorFlow 2 enables you to leverage artificial intelligence creatively and generate human-like data in the form of speech, text, images, and music. Over the next few decades, machine learning and data science will transform the finance industry. With this practical book, analysts, traders, researchers, and developers will learn how to build machine learning algorithms crucial to the

industry. You'll examine ML concepts and over 20 case studies in supervised, unsupervised, and reinforcement learning, along with natural language processing (NLP). Ideal for professionals working at hedge funds, investment and retail banks, and fintech firms, this book also delves deep into portfolio management, algorithmic trading, derivative pricing, fraud detection, asset price prediction, sentiment analysis, and chatbot development. You'll explore real-life problems faced by practitioners and learn scientifically sound solutions supported by code and examples. This book covers:

- Supervised learning regression-based models for trading strategies, derivative pricing, and portfolio management
- Supervised learning classification-based models for credit default risk prediction, fraud detection, and trading strategies
- Dimensionality reduction techniques with case studies in portfolio management, trading strategy, and yield curve construction
- Algorithms and clustering techniques for finding similar objects, with case studies in trading strategies and portfolio management
- Reinforcement learning models and techniques used for building trading strategies, derivatives hedging, and portfolio management
- NLP techniques using Python libraries such as NLTK and scikit-learn for transforming text into meaningful representations

This book, gathering the Proceedings of the 2018 Computing Conference, offers a remarkable collection of chapters covering a wide range of topics in intelligent systems, computing and their real-world applications. The Conference attracted a total of 568 submissions from pioneering researchers, scientists, industrial engineers, and students from all around the world. These submissions underwent a double-blind peer review process. Of those 568 submissions, 192 submissions (including 14 poster papers) were selected for inclusion in these proceedings. Despite computer science's comparatively brief history as a formal academic discipline, it has made a number of fundamental contributions to science and society—in fact, along with electronics, it is a founding science of the current epoch of human history ('the Information Age') and a main driver of the Information Revolution. The goal of this conference is to provide a platform for researchers to present fundamental contributions, and to be a premier venue for academic and industry practitioners to share new ideas and development experiences. This book collects state of the art chapters on all aspects of Computer Science, from classical to intelligent. It covers both the theory and applications of the latest computer technologies and methodologies. Providing the state of the art in intelligent methods and techniques for solving real-world problems, along with a vision of future research, the book will be interesting and valuable for a broad readership.

This book constitutes the proceedings of the 10th International Symposium on Cyberspace Safety and Security, CSS 2018, held in Amalfi, Italy, in October 2018. The 25 full papers presented in this volume were carefully reviewed and selected from 79 submissions. The papers focus on cybersecurity; cryptography, data security, and biometric techniques; and social security, ontologies, and smart applications.

In this book, you will learn how to create powerful machine learning based applications for a wide variety of problems leveraging different data services from the Google Cloud Platform. Finally, you will know the main difficulties that you may encounter and get appropriate strategies to overcome these difficulties and build efficient systems.

Summary Natural Language Processing in Action is your guide to creating machines that understand human language using the power of Python with its ecosystem of packages dedicated to NLP and AI. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Recent advances in deep learning empower applications to understand text and speech with extreme accuracy. The result? Chatbots that can imitate real people, meaningful resume-to-job matches, superb predictive search, and automatically generated document summaries—all at a low cost. New techniques, along with accessible tools like Keras and TensorFlow, make professional-quality NLP easier than ever before. About the Book Natural Language Processing in Action is your guide to building machines that can read and interpret human language. In it, you'll use readily available Python packages to capture the meaning in text and react accordingly. The book expands traditional NLP approaches to include neural networks, modern deep learning algorithms, and generative techniques as you tackle real-world problems like extracting dates and names, composing text, and answering free-form questions. What's inside Some sentences in this book were written by NLP!

Can you guess which ones? Working with Keras, TensorFlow, gensim, and scikit-learn Rule-based and data-based NLP Scalable pipelines About the Reader This book requires a basic understanding of deep learning and intermediate Python skills. About the Author Hobson Lane, Cole Howard, and Hannes Max Hapke are experienced NLP engineers who use these techniques in production. Table of Contents PART 1 - WORDY MACHINES Packets of thought (NLP overview) Build your vocabulary (word tokenization) Math with words (TF-IDF vectors) Finding meaning in word counts (semantic analysis) PART 2 - DEEPER LEARNING (NEURAL NETWORKS) Baby steps with neural networks (perceptrons and backpropagation) Reasoning with word vectors (Word2vec) Getting words in order with convolutional neural networks (CNNs) Loopy (recurrent) neural networks (RNNs) Improving retention with long short-term memory networks Sequence-to-sequence models and attention PART 3 - GETTING REAL (REAL-WORLD NLP CHALLENGES) Information extraction (named entity extraction and question answering) Getting chatty (dialog engines) Scaling up (optimization, parallelization, and batch processing)

Dieser Sammelband zeichnet durch zahlreiche Beispiele aus dem Unternehmensalltag ein Bild davon, vor welchen Herausforderungen Einrichtungen im Gesundheitswesen im Rahmen des Dienstleistungscontrolling stehen und wie sie diesen mit Ideen und wertvollen Handlungsempfehlungen begegnen. Untermuert wird dies durch die aktuelle, wissenschaftlich abgesicherte Theorie. Das Werk bietet eine anwendungsbezogene Darstellung ausgewählter Controllingthemen: Neben den klassischen Themen werden auch weiterführende und aktuelle Themen vorgestellt, wie die Strategiebildung, die Realisierung von Zielen und die Erfolgsmessung mittels Balanced Scorecard und Kennzahlen, aber auch die Analyse und die Bewertung von Prozessen und Leistungen in Unternehmen, um festgelegte Ziele zu erreichen. Aufgrund dieses Themenspektrums richtet sich der Sammelband einerseits an Wissenschaftler und Studenten mit den Schwerpunktfächern Controlling und Dienstleistungsmanagement, andererseits an Praktiker im

Gesundheitswesen, wie z.B. Geschäftsführer und verantwortliche Entscheidungsträger.

This open access book comprehensively covers the fundamentals of clinical data science, focusing on data collection, modelling and clinical applications. Topics covered in the first section on data collection include: data sources, data at scale (big data), data stewardship (FAIR data) and related privacy concerns. Aspects of predictive modelling using techniques such as classification, regression or clustering, and prediction model validation will be covered in the second section. The third section covers aspects of (mobile) clinical decision support systems, operational excellence and value-based healthcare. Fundamentals of Clinical Data Science is an essential resource for healthcare professionals and IT consultants intending to develop and refine their skills in personalized medicine, using solutions based on large datasets from electronic health records or telemonitoring programmes. The book's promise is "no math, no code" and will explain the topics in a style that is optimized for a healthcare audience.

If you want to build an enterprise-quality application that uses natural language text but aren't sure where to begin or what tools to use, this practical guide will help get you started. Alex Thomas, principal data scientist at Wisecube, shows software engineers and data scientists how to build scalable natural language processing (NLP) applications using deep learning and the Apache Spark NLP library. Through concrete examples, practical and theoretical explanations, and hands-on exercises for using NLP on the Spark processing framework, this book teaches you everything from basic linguistics and writing systems to sentiment analysis and search engines. You'll also explore special concerns for developing text-based applications, such as performance. In four sections, you'll learn NLP basics and building blocks before diving into application and system building: Basics: Understand the fundamentals of natural language processing, NLP on Apache Spark, and deep learning Building blocks: Learn techniques for building NLP applications—including tokenization, sentence segmentation, and named-entity recognition—and discover how and why they work Applications: Explore the design, development, and experimentation process for building your own NLP applications Building NLP systems: Consider options for productionizing and deploying NLP models, including which human languages to support This certification guide offers complete, up-to-date coverage of the PL-200 exam so you can prepare effectively. The book covers topics such as configuring Dataverse, creating Power Apps, managing processes with Power Automate, implementing chatbots, and even integrating Power Platform with other apps.

Summary Serverless Applications with Node.js walks you through building serverless apps on AWS using JavaScript. Inside, you'll discover what Claudia.js brings to the table as you build and deploy a scalable event-based serverless application, based around a pizzeria that's fully integrated with AWS services, including Lambda and API Gateway. Each chapter is filled with exercises, examples, tips, and more to make sure you're ready to bring what you've learned into your own work. Foreword by Gojko Adzic. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The benefits of cloud-hosted serverless web apps are undeniable: lower complexity, quicker time to market, and easier scalability than traditional, server-dependent designs. And thanks to JavaScript support in AWS Lambda and powerful new serverless API tools like the Claudia.js library, you can build and deploy serverless apps end to end without learning a new language. About the Book Serverless Applications with Node.js teaches you to design and build serverless web apps on AWS using JavaScript, Node, and Claudia.js. You'll master the basics of writing AWS Lambda functions, along with core serverless patterns like API Gateway. Along the way, you'll practice your new skills by building a working chatbot and a voice assistant with Amazon Alexa. You'll also discover techniques for migrating existing apps to a serverless platform. What's inside Authentication and database storage Asynchronous functions Interesting real-world examples Developing serverless microservices About the Reader For web developers comfortable with JavaScript and Node.js. About the Author Slobodan Stojanovi? and Aleksandar Simovi? are AWS Serverless Heroes and core contributors to the Claudia.js project. They are also coauthors of Desole, an open source serverless errortracking tool, and the lead developers of Claudia Bot Builder. Table of Contents PART 1 - Serverless pizzeria Introduction to serverless with Claudia Building your first serverless API Asynchronous work is easy, we Promise() Pizza delivery: Connecting an external service Houston, we have a problem! Level up your API Working with files PART 2 - Let's talk When pizza is one message away: Chatbots Typing... Async and delayed responses Jarvis, I mean Alexa, order me a pizza Paying for pizza Migrating to serverless Real-world case studies appendix A - Installation and configuration appendix B - Facebook Messenger, Twilio, and Alexa configuration appendix C - Stripe and MongoDB setup appendix D - The pizza recipe Create powerful, useful, fast, and fun chat bots that make Slack better About This Book This is the first developer's guide to programming for Slack It covers everything you need to create chat bots for Slack's conversational UI It's full of realistic examples, usable code, and lucid explanations on everything you need to know Who This Book Is For This is a book for software developers who want to build Slack bots for their own company's use or for customers. You need to know JavaScript. What You Will Learn Receive and send messages sent in Slack Respond to user commands Process natural language Perform useful tasks on commands such as fetching data from external sources Use webhooks and slash commands Create a Google Classroom and add customized information for each individual class In Detail Slack promises that its users will "be less busy." Slack bots interact with users in Slack chatrooms, providing useful immediate information, and automating work. This book gives you everything you need to build powerful and useful Slack bots. You'll see how to hook into the Slack API to create software that can read and post to chatrooms, respond to commands and hints given in natural conversational language, and build fun and useful bots for your own place of work, both as a front end to your own service and to distribute and share as apps. You can even sell your bots and build a business as a Slack bot developer. Throughout the book, you'll build useful and fun example applications that you can modify for your own situations. These range from simple, fun applications to liven up discussions to useful, data-driven apps to help you make decisions quickly and manage work. Style and approach This is a friendly step-by-step guide to building powerful, friendly Slack bots.

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

A hands-on, application-based introduction to machine learning and artificial intelligence (AI) that guides young readers through creating compelling AI-powered games and applications using the Scratch programming language. Machine learning (also known as ML) is one of the building blocks of AI, or artificial intelligence. AI is based on the idea that computers can learn on their own, with your help. Machine Learning for Kids will introduce you to machine learning, painlessly. With this book and its free, Scratch-based, award-winning companion website, you'll see how easy it is to add machine learning to your own projects. You don't even need to know how to code! As you work through the book you'll discover how machine learning systems can be taught to recognize text, images, numbers, and sounds, and how to train your models to improve their accuracy. You'll turn your models into fun computer games and apps, and see what happens when they get confused by bad data. You'll build 13 projects step-by-step from the ground up, including:

- Rock, Paper, Scissors game that recognizes your hand shapes
- An app that recommends movies based on other movies that you like
- A computer character that reacts to insults and compliments
- An interactive virtual assistant (like Siri or Alexa) that obeys commands
- An AI version of Pac-Man, with a smart character that knows how to avoid ghosts

NOTE: This book includes a Scratch tutorial for beginners, and step-by-step instructions for every project. Ages 12+

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Artificial Intelligence(AI), gets your system to think smart and intelligent. This book is packed with some of the smartest and easy-peasy examples through which you will learn the fundamentals of AI. You will have acquired the foundation of AI and understood the practical case studies in this book.

This book offers a highly accessible introduction to natural language processing, the field that supports a variety of language technologies, from predictive text and email filtering to automatic summarization and translation. With it, you'll learn how to write Python programs that work with large collections of unstructured text. You'll access richly annotated datasets using a comprehensive range of linguistic data structures, and you'll understand the main algorithms for analyzing the content and structure of written communication. Packed with examples and exercises, Natural Language Processing with Python will help you:

- Extract information from unstructured text, either to guess the topic or identify "named entities"
- Analyze linguistic structure in text, including parsing and semantic analysis
- Access popular linguistic databases, including WordNet and treebanks
- Integrate techniques drawn from fields as diverse as linguistics and artificial intelligence

This book will help you gain practical skills in natural language processing using the Python programming language and the Natural Language Toolkit (NLTK) open source library. If you're interested in developing web applications, analyzing multilingual news sources, or documenting endangered languages -- or if you're simply curious to have a programmer's perspective on how human language works -- you'll find Natural Language Processing with Python both fascinating and immensely useful.

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

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

- This is the latest practice test to pass the PL-900 Microsoft Power Platform Fundamentals Exam. - It contains 160 Questions and Answers. - All the questions are 100% valid and stable. - You can reply on this practice test to pass the exam with a good mark and in the first attempt.

Cut through the noise and get real results with a step-by-step approach to learning Python 3.X programming Key Features Ideal for the Python beginner who is getting started for the first time A step-by-step Python tutorial with exercises and activities that help build key skills Structured to let you progress at your own pace, on your own terms Use your physical print copy to redeem free access to the online interactive edition Book Description You already know you want to learn Python, and a smarter way to learn Python 3 is to learn by doing. The Python Workshop focuses on building up your practical skills so that you can work towards building up your machine learning skills as a data scientist, write scripts that help automate your life and save you time, or even create your own games and desktop applications. You'll learn from real examples that lead to real results. Throughout The Python Workshop, you'll take an engaging step-by-step approach to understanding Python. You won't have to sit through any unnecessary theory. If you're short on time you can jump into a single exercise each day or spend an entire weekend learning about Python scripting. It's your choice. Learning on your terms, you'll build up and reinforce key skills in a way that feels rewarding. Every physical print copy of The Python Workshop unlocks access to the interactive edition. With videos detailing all exercises and activities, you'll always have a guided solution. You can also benchmark yourself against assessments, track progress, and receive free content updates. You'll even earn a secure credential that you can share and verify online upon completion. It's a premium learning experience that's included with your printed copy. To redeem, follow the instructions located at the start of your Python book. Fast-paced and direct, The Python Workshop is the ideal companion for Python beginners. You'll build and iterate on your code like a software developer, learning along the way. This process means that you'll find that your new skills stick, embedded as best practice. A solid foundation for the years ahead. What you will learn Learn how to write clean and concise code with Python 3 Understand classes and object-oriented programming Tackle entry-level data science and create engaging visualizations Use Python to create responsive, modern web applications Automate essential day-to-day tasks with Python scripts Get started with predictive Python machine learning Who this book is for This book is designed for professionals, students, and hobbyists who want to learn Python and apply it to solve challenging real-world problems. Although this is a beginner's book, it will help if you already know standard programming topics, such as variables, if-else statements, and functions. Experience with another object-oriented program is beneficial, but not mandatory.

Creative Coding in Python presents over 30 creative projects that teach kids how to code in the easy and intuitive programming language, Python. Creative Coding in Python teaches the fundamentals of computer programming and demonstrates how to code 30+ fun, creative projects using Python, a free, intuitive, open-source programming language that's one of the top five most popular worldwide and one of the most popular Google search terms in the U.S. Computer science educator Sheena Vaidyanathan helps kids understand the fundamental ideas of computer programming and the process of computational thinking using illustrations, flowcharts, and pseudocode, then shows how to apply those essentials to code exciting projects in Python: Chatbots: Discover variables, strings, integers, and more to design conversational programs. Geometric art: Use turtle graphics to create original masterpieces. Interactive fiction: Explore booleans and conditionals to invent "create your own adventure" games. Dice games: Reuse code to devise games of chance. Arcade games and apps: Understand GUI (graphical user interfaces) and create your own arcade games and apps. What's next? Look at exciting ways to use your powerful new skills and expand your knowledge of coding in Python. Creative Coding in Python gives kids the tools they need to create their own computer programs. Eat, drink, and be murdered. Someone is poisoning diners in New Orleans' best restaurants. Now it's up to chef and homicide cop Caleb Rooney to catch a killer--who has an appetite for revenge. BookShots LIGHTNING-FAST STORIES BY JAMES

PATTERSON Novels you can devour in a few hours Impossible to stop reading All original content from James Patterson Make NLP easy by building chatbots and models, and executing various NLP tasks to gain data-driven insights from raw text data Key Features Get familiar with key natural language processing (NLP) concepts and terminology Explore the functionalities and features of popular NLP tools Learn how to use Python programming and third-party libraries to perform NLP tasks Book Description Do you want to learn how to communicate with computer systems using Natural Language Processing (NLP) techniques, or make a machine understand human sentiments? Do you want to build applications like Siri, Alexa, or chatbots, even if you've never done it before? With The Natural Language Processing Workshop, you can expect to make consistent progress as a beginner, and get up to speed in an interactive way, with the help of hands-on activities and fun exercises. The book starts with an introduction to NLP. You'll study different approaches to NLP tasks, and perform exercises in Python to understand the process of preparing datasets for NLP models. Next, you'll use advanced NLP algorithms and visualization techniques to collect datasets from open websites, and to summarize and generate random text from a document. In the final chapters, you'll use NLP to create a chatbot that detects positive or negative sentiment in text documents such as movie reviews. By the end of this book, you'll be equipped with the essential NLP tools and techniques you need to solve common business problems that involve processing text. What you will learn Obtain, verify, clean and transform text data into a correct format for use Use methods such as tokenization and stemming for text extraction Develop a classifier to classify comments in Wikipedia articles Collect data

from open websites with the help of web scraping Train a model to detect topics in a set of documents using topic modeling Discover techniques to represent text as word and document vectors Who this book is for This book is for beginner to mid-level data scientists, machine learning developers, and NLP enthusiasts. A basic understanding of machine learning and NLP is required to help you grasp the topics in this workshop more quickly.

New edition of the bestselling guide to deep reinforcement learning and how it's used to solve complex real-world problems. Revised and expanded to include multi-agent methods, discrete optimization, RL in robotics, advanced exploration techniques, and more Key Features Second edition of the bestselling introduction to deep reinforcement learning, expanded with six new chapters Learn advanced exploration techniques including noisy networks, pseudo-count, and network distillation methods Apply RL methods to cheap hardware robotics platforms Book Description Deep Reinforcement Learning Hands-On, Second Edition is an updated and expanded version of the bestselling guide to the very latest reinforcement learning (RL) tools and techniques. It provides you with an introduction to the fundamentals of RL, along with the hands-on ability to code intelligent learning agents to perform a range of practical tasks. With six new chapters devoted to a variety of up-to-the-minute developments in RL, including discrete optimization (solving the Rubik's Cube), multi-agent methods, Microsoft's TextWorld environment, advanced exploration techniques, and more, you will come away from this book with a deep understanding of the latest innovations in this emerging field. In addition, you will gain actionable insights into such topic areas as deep Q-networks, policy gradient methods, continuous control problems, and highly scalable, non-gradient methods. You will also discover how to build a real hardware robot trained with RL for less than \$100 and solve the Pong environment in just 30 minutes of training using step-by-step code optimization. In short, Deep Reinforcement Learning Hands-On, Second Edition, is your companion to navigating the exciting complexities of RL as it helps you attain experience and knowledge through real-world examples. What you will learn Understand the deep learning context of RL and implement complex deep learning models Evaluate RL methods including cross-entropy, DQN, actor-critic, TRPO, PPO, DDPG, D4PG, and others Build a practical hardware robot trained with RL methods for less than \$100 Discover Microsoft's TextWorld environment, which is an interactive fiction games platform Use discrete optimization in RL to solve a Rubik's Cube Teach your agent to play Connect 4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI chatbots Discover advanced exploration techniques, including noisy networks and network distillation techniques Who this book is for Some fluency in Python is assumed. Sound understanding of the fundamentals of deep learning will be helpful. This book is an introduction to deep RL and requires no background in RL

[Copyright: 7d28ea4ba7d787cd1d53c0a5421e0e26](#)