

| | |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Record Nr. | UNINA9910818902203321 |
| Autore | Rothman Denis |
| Titolo | Artificial Intelligence by Example : Acquire Advanced AI, Machine Learning, and Deep Learning Design Skills |
| Pubbl/distr/stampa | Birmingham : , : Packt Publishing, Limited, , 2020 ©2020 |
| ISBN | 1-83921-281-0 |
| Edizione | [2nd ed.] |
| Descrizione fisica | 1 online resource (579 pages) |
| Disciplina | 6.3 |
| Soggetti | Artificial intelligence |
| Lingua di pubblicazione | Inglese |
| Formato | Materiale a stampa |
| Livello bibliografico | Monografia |
| Note generali | Previous edition published: 2018. |
| Nota di bibliografia | Includes bibliographical references and index. |
| Nota di contenuto | Cover -- Copyright -- Packt Page -- Contributors -- Table of Contents -- Preface -- Chapter 1: Getting Started with Next-Generation Artificial Intelligence through Reinforcement Learning -- Reinforcement learning concepts -- How to adapt to machine thinking and become an adaptive thinker -- Overcoming real-life issues using the three-step approach -- Step 1 - describing a problem to solve: MDP in natural language -- Watching the MDP agent at work -- Step 2 - building a mathematical model: the mathematical representation of the Bellman equation and MDP -- From MDP to the Bellman equation -- Step 3 - writing source code: implementing the solution in Python -- The lessons of reinforcement learning -- How to use the outputs -- Possible use cases -- Machine learning versus traditional applications -- Summary -- Questions -- Further reading -- Chapter 2: Building a Reward Matrix - Designing Your Datasets -- Designing datasets - where the dream stops and the hard work begins -- Designing datasets -- Using the McCulloch-Pitts neuron -- The McCulloch-Pitts neuron -- The Python-TensorFlow architecture -- Logistic activation functions and classifiers -- Overall architecture -- Logistic classifier -- Logistic function -- Softmax -- Summary -- Questions -- Further reading -- Chapter 3: Machine Intelligence - Evaluation Functions and Numerical Convergence -- Tracking down what to measure and deciding how to measure it -- Convergence -- Implicit convergence -- Numerically controlled gradient descent convergence -- Evaluating beyond human |

analytic capacity -- Using supervised learning to evaluate a result that surpasses human analytic capacity -- Summary -- Questions -- Further reading -- Chapter 4: Optimizing Your Solutions with K-Means Clustering -- Dataset optimization and control -- Designing a dataset and choosing an ML/DL model. Approval of the design matrix -- Implementing a k-means clustering solution -- The vision -- The data -- The strategy -- The k-means clustering program -- The mathematical definition of k-means clustering -- The Python program -- Saving and loading the model -- Analyzing the results -- Bot virtual clusters as a solution -- The limits of the implementation of the k-means clustering algorithm -- Summary -- Questions -- Further reading -- Chapter 5: How to Use Decision Trees to Enhance K-Means Clustering -- Unsupervised learning with KMC with large datasets -- Identifying the difficulty of the problem -- NP-hard - the meaning of P -- NP-hard - the meaning of non-deterministic -- Implementing random sampling with mini-batches -- Using the LLN -- The CLT -- Using a Monte Carlo estimator -- Trying to train the full training dataset -- Training a random sample of the training dataset -- Shuffling as another way to perform random sampling -- Chaining supervised learning to verify unsupervised learning -- Preprocessing raw data -- A pipeline of scripts and ML algorithms -- Step 1 - training and exporting data from an unsupervised ML algorithm -- Step 2 - training a decision tree -- Step 3 - a continuous cycle of KMC chained to a decision tree -- Random forests as an alternative to decision trees -- Summary -- Questions -- Further reading -- Chapter 6: Innovating AI with Google Translate -- Understanding innovation and disruption in AI -- Is AI disruptive? -- AI is based on mathematical theories that are not new -- Neural networks are not new -- Looking at disruption - the factors that are making AI disruptive -- Cloud server power, data volumes, and web sharing of the early 21st century -- Public awareness -- Inventions versus innovations -- Revolutionary versus disruptive solutions -- Where to start? -- Discover a world of opportunities with Google Translate. Getting started -- The program -- The header -- Implementing Google's translation service -- Google Translate from a linguist's perspective -- Playing with the tool -- Linguistic assessment of Google Translate -- AI as a new frontier -- Lexical field and polysemy -- Exploring the frontier - customizing Google Translate with a Python program -- k-nearest neighbor algorithm -- Implementing the KNN algorithm -- The knn_polysemy.py program -- Implementing the KNN function in Google_Translate_Customized.py -- Conclusions on the Google Translate customized experiment -- The disruptive revolutionary loop -- Summary -- Questions -- Further reading -- Chapter 7: Optimizing Blockchains with Naive Bayes -- Part I - the background to blockchain technology -- Mining bitcoins -- Using cryptocurrency -- PART II - using blockchains to share information in a supply chain -- Using blockchains in the supply chain network -- Creating a block -- Exploring the blocks -- Part III - optimizing a supply chain with naive Bayes in a blockchain process -- A naive Bayes example -- The blockchain anticipation novelty -- The goal - optimizing storage levels using blockchain data -- Implementation of naive Bayes in Python -- Gaussian naive Bayes -- Summary -- Questions -- Further reading -- Chapter 8: Solving the XOR Problem with a Feedforward Neural Network -- The original perceptron could not solve the XOR function -- XOR and linearly separable models -- Linearly separable models -- The XOR limit of a linear model, such as the original perceptron -- Building an FNN from scratch -- Step 1 - defining an FNN -- Step 2 - an example of how two children can solve

the XOR problem every day -- Implementing a vintage XOR solution in Python with an FNN and backpropagation -- A simplified version of a cost function and gradient descent -- Linear separability was achieved. Applying the FNN XOR function to optimizing subsets of data -- Summary -- Questions -- Further reading -- Chapter 9: Abstract Image Classification with Convolutional Neural Networks (CNNs) -- Introducing CNNs -- Defining a CNN -- Initializing the CNN -- Adding a 2D convolution layer -- Kernel -- Shape -- ReLU -- Pooling -- Next convolution and pooling layer -- Flattening -- Dense layers -- Dense activation functions -- Training a CNN model -- The goal -- Compiling the model -- The loss function -- The Adam optimizer -- Metrics -- The training dataset -- Data augmentation -- Loading the data -- The testing dataset -- Data augmentation on the testing dataset -- Loading the data -- Training with the classifier -- Saving the model -- Next steps -- Summary -- Questions -- Further reading and references -- Chapter 10: Conceptual Representation Learning -- Generating profit with transfer learning -- The motivation behind transfer learning -- Inductive thinking -- Inductive abstraction -- The problem AI needs to solve -- The gap concept -- Loading the trained TensorFlow 2.x model -- Loading and displaying the model -- Loading the model to use it -- Defining a strategy -- Making the model profitable by using it for another problem -- Domain learning -- How to use the programs -- The trained models used in this section -- The trained model program -- Gap - loaded or underloaded -- Gap - jammed or open lanes -- Gap datasets and subsets -- Generalizing the (the gap conceptual dataset) -- The motivation of conceptual representation learning metamodels applied to dimensionality -- The curse of dimensionality -- The blessing of dimensionality -- Summary -- Questions -- Further reading -- Chapter 11: Combining Reinforcement Learning and Deep Learning -- Planning and scheduling today and tomorrow -- A real-time manufacturing process. Amazon must expand its services to face competition -- A real-time manufacturing revolution -- CRLMM applied to an automated apparel manufacturing process -- An apparel manufacturing process -- Training the CRLMM -- Generalizing the unit training dataset -- Food conveyor belt processing - positive p and negative n gaps -- Running a prediction program -- Building the RL-DL-CRLMM -- A circular process -- Implementing a CNN-CRLMM to detect gaps and optimize -- Q-learning - MDP -- MDP inputs and outputs -- The optimizer -- The optimizer as a regulator -- Finding the main target for the MDP function -- A circular model - a stream-like system that never starts nor ends -- Summary -- Questions -- Further reading -- Chapter 12: AI and the Internet of Things (IoT) -- The public service project -- Setting up the RL-DL-CRLMM model -- Applying the model of the CRLMM -- The dataset -- Using the trained model -- Adding an SVM function -- Motivation - using an SVM to increase safety levels -- Definition of a support vector machine -- Python function -- Running the CRLMM -- Finding a parking space -- Deciding how to get to the parking lot -- Support vector machine -- The itinerary graph -- The weight vector -- Summary -- Questions -- Further reading -- Chapter 13: Visualizing Networks with TensorFlow 2.x and TensorBoard -- Exploring the output of the layers of a CNN in two steps with TensorFlow -- Building the layers of a CNN -- Processing the visual output of the layers of a CNN -- Analyzing the visual output of the layers of a CNN -- Analyzing the accuracy of a CNN using TensorBoard -- Getting started with Google Colaboratory -- Defining and training the model -- Introducing some of the measurements -- Summary -- Questions -- Further reading.

Chapter 14: Preparing the Input of Chatbots with Restricted Boltzmann Machines (RBMs) and Principal Component Analysis (PCA).

Sommario/riassunto

Artificial Intelligence (AI) gets your system to think smart and learn intelligently. This book is packed with some of the smartest trending examples with which you will learn the fundamentals of AI. By the end, you will have acquired the basics of AI by practically applying the examples in this book.
