

1. Record Nr.	UNINA9910584485003321
Autore	Fergus Paul
Titolo	Applied deep learning : tools, techniques, and implementation // Paul Fergus, Carl Chalmers
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2022] ©2022
ISBN	3-031-04420-7
Descrizione fisica	1 online resource (355 pages)
Collana	Computational Intelligence Methods and Applications
Disciplina	006.31
Soggetti	Deep learning (Machine learning)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Acknowledgements -- Contents -- List of Figures -- List of Tables -- Part I: Introduction and Overview -- Chapter 1: Introduction -- 1.1 Artificial Intelligence, Machine Learning, Deep Learning -- 1.1.1 Artificial Intelligence -- 1.1.2 Machine Learning -- 1.1.3 Deep Learning -- 1.1.4 How they Come Together -- 1.2 Artificial Intelligence Is Driving Innovation -- 1.2.1 Transforming Healthcare -- 1.2.2 Protecting Wildlife -- 1.2.3 Securing the Environment -- 1.3 Tools, Frameworks and Hardware -- 1.3.1 Building Intelligent Applications -- 1.3.2 Python, Notebooks and Environments -- 1.3.3 Pre-Processing -- 1.3.4 Machine Learning -- 1.3.5 Deep Learning -- 1.3.6 Inferencing -- 1.4 How this Book Is Organised -- 1.5 Who Should Read this Book -- 1.6 Summary -- References -- Part II: Foundations of Machine Learning -- Chapter 2: Fundamentals of Machine Learning -- 2.1 What Is Machine Learning? -- 2.1.1 Formal and Non-Formal Definition -- 2.1.2 How AI and Machine Learning Differs from Conventional Software Development -- 2.1.2.1 Rewriting the Rules -- 2.1.2.2 Intelligent Decision Making -- 2.2 Machine Learning Tribes -- 2.2.1 Connectionists -- 2.2.2 Evolutionists -- 2.2.3 Bayesians -- 2.2.4 Symbolists -- 2.2.5 Analogists -- 2.3 Data Management -- 2.3.1 Data Types and Data Objects -- 2.3.1.1 Numerical -- 2.3.1.2 Textual -- 2.3.1.3 Categorical -- 2.3.1.4 Timeseries -- 2.3.2 Data Structure -- 2.3.2.1 Data Objects -- 2.3.3 Datasets -- 2.3.4 Exploratory Data Analysis -- 2.3.4.1 What Is Exploratory Data Analysis -- 2.3.4.2 Data

Distributions -- 2.3.4.3 Validate Assumptions -- 2.3.4.4 Feature Engineering -- 2.4 Learning Problems -- 2.4.1 Supervised Machine Learning -- 2.4.2 Semi-Supervised Machine Learning -- 2.4.3 Un-Supervised Machine Learning -- 2.4.4 Regression -- 2.4.5 Reinforcement Learning -- 2.5 Evaluating Machine Learning Models. 2.6 Summary -- References -- Chapter 3: Supervised Learning -- 3.1 Basic Concepts -- 3.2 Supervised Learning Tasks -- 3.2.1 Data Extraction -- 3.2.2 Data Preparation -- 3.2.2.1 Data Size -- 3.2.2.2 Missing Data -- 3.2.2.3 Textual Data -- One Hot Encoding -- 3.2.2.4 Value Ranges (Normalisation and Scaling) -- Scaling -- Normalisation -- Standardisation -- 3.2.2.5 Distribution -- 3.2.2.6 Class Balance -- 3.2.2.7 Correlation Between Features -- 3.2.3 Feature Engineering -- 3.2.3.1 Feature Selection -- 3.2.3.2 Dimensionality Reduction -- 3.2.4 Selecting a Training Algorithm -- 3.3 Supervised Algorithms -- 3.3.1 Linear Regression -- 3.3.2 Logistic Regression -- 3.3.3 Linear Discriminate Analysis -- 3.3.4 Support Vector Machine -- 3.3.5 Random Forest -- 3.3.6 Naive Bayes -- 3.3.7 K-Nearest Neighbours -- 3.4 Summary -- References -- Chapter 4: Un-Supervised Learning -- 4.1 Basic Concepts -- 4.2 Clustering -- 4.2.1 Hierarchical Clustering -- 4.2.2 K-Means -- 4.2.3 Mixture Models -- 4.2.4 DBSCAN -- 4.2.5 Optics Algorithm -- 4.3 Principal Component Analysis -- 4.4 Association Rule Mining -- 4.5 Summary -- References -- Chapter 5: Performance Evaluation Metrics -- 5.1 Introduction to Model Evaluation -- 5.1.1 Evaluation Challenges -- 5.1.2 Taxonomy of Classifier Evaluation Metrics -- 5.2 Classification Accuracy -- 5.3 Train, Test and Validation Sets -- 5.4 Underfitting and Overfitting -- 5.5 Supervised Learning Evaluation Metrics -- 5.5.1 Confusion Matrix -- 5.5.1.1 Accuracy -- 5.5.1.2 Precision -- 5.5.1.3 Recall (Sensitivity) -- 5.5.1.4 Specificity -- 5.5.1.5 False Positive Rate -- 5.5.1.6 F1-Score -- 5.5.2 Receiver Operating Characteristic -- 5.5.3 Regression Metrics -- 5.5.3.1 Mean Square Error (MSE) -- 5.5.3.2 MAE -- 5.5.3.3 R2 (Coefficient of Determination) -- 5.6 Probability Scoring Methods -- 5.6.1 Log Loss Score -- 5.6.2 Brier Score. 5.7 Cross-Validation -- 5.7.1 Challenge of Evaluating Classifiers -- 5.7.2 K-Fold Cross-Validation -- 5.8 Un-Supervised Learning Evaluation Metric -- 5.8.1 Elbow Method -- 5.8.2 Davies-Bouldin Index -- 5.8.3 Dunn Index -- 5.8.4 Silhouette Coefficient -- 5.9 Summary -- References -- Part III: Deep Learning Concepts and Techniques -- Chapter 6: Introduction to Deep Learning -- 6.1 So what's the Difference Between DL and ML? -- 6.2 Introduction to Deep Learning -- 6.3 Artificial Neural Networks -- 6.3.1 Perceptrons -- 6.3.2 Neural Networks -- 6.3.3 Activation Functions -- 6.3.4 Multi-Class Classification Considerations -- 6.3.5 Cost Functions and Optimisers -- 6.3.6 Backpropagation -- 6.3.7 The Vanishing Gradient -- 6.3.8 Weight Initialisation -- 6.3.9 Regularisation -- 6.4 Convolutional Neural Networks -- 6.4.1 Image Filters and Kernels -- 6.4.2 Convolutional Layers -- 6.4.3 Pooling Layers -- 6.4.4 Transfer Learning -- 6.5 Summary -- References -- Chapter 7: Image Classification and Object Detection -- 7.1 Hardware Accelerated Deep Learning -- 7.1.1 Training and Associated Hardware -- 7.1.1.1 Development Systems -- 7.1.1.2 Training Systems -- 7.1.1.3 Inferencing Systems -- 7.1.2 Tensor Processing Unit (TPU) -- 7.1.3 Other Hardware Considerations -- 7.2 Object Recognition -- 7.2.1 Image Classification -- 7.2.2 Object Detection -- 7.2.3 Semantic Segmentation -- 7.2.4 Object Segmentation -- 7.3 Model Architectures -- 7.3.1 Single Shot Detector (SSD) -- 7.3.2 YOLO Family -- 7.3.3 R-CNN -- 7.3.4 Fast-RCNN -- 7.3.5 Faster-RCNN -- 7.3.6 EfficientNet -- 7.3.7 Comparing Architectures -- 7.3.7.1 Key Findings -- 7.3.7.2 Most Accurate --

7.3.7.3 Fastest -- 7.4 Evaluation Metrics -- 7.4.1 Confidence Score -- 7.4.2 Intersection over Union -- 7.4.3 Mean Average Precision (mAP) -- 7.5 Summary -- References.

Chapter 8: Deep Learning Techniques for Time Series Modelling -- 8.1 Introduction to Time-Series Data -- 8.2 Recurrent Neural Network -- 8.2.1 Developing RNNs for Time Series Forecasting -- 8.3 Long-Term Short-Term Memory -- 8.4 Gated Recurrent Unit -- 8.5 One Dimensional Convolutional Neural Network -- 8.6 Summary -- References -- Chapter 9: Natural Language Processing -- 9.1 Introduction to Natural Language Processing -- 9.1.1 Tokenisation -- 9.1.2 Stemming -- 9.1.3 Lemmatization -- 9.1.4 Stop Words -- 9.1.5 Phrase Matching and Vocabulary -- 9.2 Text Classification -- 9.2.1 Text Feature Extraction -- 9.3 Sentiment Analysis -- 9.4 Topic Modelling -- 9.4.1 Latent Semantic Analysis (LSA) -- 9.4.2 Latent Dirichlet Allocation -- 9.4.3 Non-negative Matrix Factorization -- 9.5 Deep Learning for NLP -- 9.5.1 Word Embeddings -- 9.5.2 Word Embedding Algorithms -- 9.5.2.1 Embedding Layer -- 9.5.2.2 Word2Vec -- 9.5.2.3 GloVe -- 9.5.2.4 Natural Language Understanding and Generation -- 9.6 Real-World Applications -- 9.6.1 Chat Bots -- 9.6.2 Smart Speakers -- 9.7 Summary -- References -- Chapter 10: Deep Generative Models -- 10.1 Autoencoders -- 10.1.1 Autoencoder Basics -- 10.1.2 Autoencoder for Dimensionality Reduction -- 10.1.3 Autoencoder for Images -- 10.1.4 Stacked Autoencoders -- 10.1.5 Generative Adversarial Networks (GANs) -- 10.1.5.1 GANs Network Architectures -- 10.2 Summary -- References -- Chapter 11: Deep Reinforcement Learning -- 11.1 What Is Reinforcement Learning? -- 11.2 Reinforcement Learning Definitions -- 11.3 Domain Selection for Reinforcement Learning -- 11.4 State-Action Pairs and Complex Probability Distributions of Reward -- 11.5 Neural Networks and Reinforcement Learning -- 11.6 The Deep Reinforcement Learning Process -- 11.7 Practical Applications of Deep Reinforcement Learning -- 11.8 Summary -- References.

Part IV: Enterprise Machine Learning -- Chapter 12: Accelerated Machine Learning -- 12.1 Introduction -- 12.1.1 CPU/GPU Based Clusters -- 12.2 CPU Accelerated Computing -- 12.2.1 Distributed Accelerated Computing Frameworks -- 12.2.1.1 Local Vs Distributed -- 12.2.1.2 Benefits of Scaling Out -- 12.2.1.3 Hadoop -- 12.2.1.4 Apache Spark -- 12.3 Introduction to DASK -- 12.3.1 DASK Arrays -- 12.3.2 Scikit Learn and DASK Integration (DASK ML) -- 12.3.3 Scikit Learn Joblib -- 12.4 GPU Computing -- 12.4.1 Introduction to GPU Hardware -- 12.4.2 Introduction to NVIDIA Accelerated Computing -- 12.4.3 CUDA -- 12.4.4 CUDA Accelerated Computing Libraries -- 12.5 RAPIDS -- 12.5.1 cuDF Analytics -- 12.5.2 cuML Machine Learning -- 12.5.3 cuGraph Graph Analytics -- 12.5.4 Apache Arrow -- 12.6 Summary -- References -- Chapter 13: Deploying and Hosting Machine Learning Models -- 13.1 Introduction to Deployment -- 13.1.1 Why Is Model Deployment Important -- 13.1.2 Enabling MLOps -- 13.1.3 MLOps Frameworks -- 13.1.4 MLOps Application Programmable Interfaces API's -- 13.2 Preparing a Model -- 13.2.1 Model Formats -- 13.2.1.1 ProtoBuf (pb) -- 13.2.1.2 ONNX (.ONNX) -- 13.2.1.3 Keras h5 (.h5) -- 13.2.1.4 TensorFlow SavedModel Format -- 13.2.1.5 Scikit-Learn (.pkl) -- 13.2.1.6 IOS Platform (.mlmodel) -- 13.2.1.7 PyTorch (.pt) -- 13.2.2 Freezing and Exporting Models -- 13.2.3 Model Optimisation -- 13.2.4 Deploying the TFLite Model and Undertaking Inference -- 13.3 Web Deployment -- 13.3.1 Flask -- 13.3.2 Why Use Flask -- 13.3.3 Working and Developing in Flask -- 13.4 Summary -- References -- Chapter 14: Enterprise Machine Learning Serving -- 14.1 Docker -- 14.1.1 What Is Docker -- 14.1.2 Working with Docker --

14.1.2.1 Using Docker -- 14.1.2.2 What's a Container -- 14.1.2.3
Docker Run -- 14.1.2.4 Container Lifecycle -- 14.1.2.5 Building
Custom Dockers -- 14.1.3 Docker Compose.
14.1.4 Docker Volume and Mount.
