

# Deep Learning With C Net And Kelp Net The Ultimat

When somebody should go to the books stores, search start by shop, shelf by shelf, it is truly problematic. This is why we give the book compilations in this website. It will very ease you to see guide **Deep Learning With C Net And Kelp Net The Ultimat** as you such as.

By searching the title, publisher, or authors of guide you in point of fact want, you can discover them rapidly. In the house, workplace, or perhaps in your method can be all best area within net connections. If you intention to download and install the Deep Learning With C Net And Kelp Net The Ultimat, it is categorically simple then, in the past currently we extend the link to purchase and create bargains to download and install Deep Learning With C Net And Kelp Net The Ultimat consequently simple!

*Deep Learning With C Net And Kelp  
Net The Ultimat*

2019-11-12

**PATEL AMY**

Deep Belief Nets in C++ and CUDA C: Volume 1 BALIGE PUBLISHING

In this comprehensive project focusing on Hepatitis C classification and prediction, the journey begins with a meticulous exploration of the dataset. Through Python, Scikit-Learn, Keras, and TensorFlow, the project aims to develop an effective model to predict Hepatitis C based on given features. The dataset's attributes are systematically examined, and their distributions are analyzed to uncover insights into potential correlations and patterns. The subsequent step involves categorizing the feature distributions. This phase sheds light on the underlying characteristics of each attribute, facilitating the understanding of their roles in influencing the target variable. This categorization lays the foundation for feature scaling and preprocessing, ensuring that the data is optimized for machine learning. The core of the project revolves around the development of machine learning models. Employing Scikit-Learn, various classification algorithms are applied, including K-Nearest Neighbors (KNN), Decision Trees, Random Forests, Naive Bayes, Gradient Boosting, AdaBoost, Light Gradient Boosting, Multi-Layer Perceptron, and XGBoost. The models are fine-tuned using Grid Search to optimize hyperparameters, enhancing their performance and generalization capability. Taking the project a step further, deep learning techniques are harnessed to tackle the Hepatitis C classification challenge. A key component is the construction of an Artificial Neural Network (ANN) using Keras and TensorFlow. This ANN leverages layers of interconnected nodes to learn

complex patterns within the data. LSTM, FNN, RNN, DBN, and Autoencoders are also explored, offering a comprehensive understanding of deep learning's versatility. To evaluate the models' performances, an array of metrics are meticulously employed. Metrics such as accuracy, precision, recall, F1-score, and AUC-ROC are meticulously calculated. The significance of each metric is meticulously explained, underpinning the assessment of a model's true predictive power and its potential weaknesses. The evaluation phase emerges as a pivotal aspect, accentuated by an array of comprehensive metrics. Performance assessment encompasses metrics such as accuracy, precision, recall, F1-score, and ROC-AUC. Cross-validation and learning curves are strategically employed to mitigate overfitting and ensure model generalization. Furthermore, visual aids such as ROC curves and confusion matrices provide a lucid depiction of the models' interplay between sensitivity and specificity. The culmination of the project involves the creation of a user-friendly Graphical User Interface (GUI) using PyQt. The GUI enables users to interact seamlessly with the models, facilitating data input, model selection, and prediction execution. A detailed description of the GUI's components, including buttons, checkboxes, and interactive plots, highlights its role in simplifying the entire classification process. In a comprehensive journey of exploration, experimentation, and analysis, this project effectively marries data science and machine learning. By thoroughly examining the dataset, engineering features, utilizing a diverse range of machine learning models, harnessing the capabilities of deep learning, evaluating performance metrics, and creating an intuitive GUI, the project encapsulates the multi-faceted nature of modern data-driven endeavors.

**Data Science For Programmer: A Project-Based Approach**

**With Python GUI** BALIGE PUBLISHING

Create and unleash the power of neural networks by implementing C# and .Net code Key FeaturesGet a strong foundation of neural networks with access to various machine learning and deep learning librariesReal-world case studies illustrating various neural network techniques and architectures used by practitionersCutting-edge coverage of Deep Networks, optimization algorithms, convolutional networks, autoencoders and many moreBook Description Neural networks have made a surprise comeback in the last few years and have brought tremendous innovation in the world of artificial intelligence. The goal of this book is to provide C# programmers with practical guidance in solving complex computational challenges using neural networks and C# libraries such as CNTK, and TensorFlowSharp. This book will take you on a step-by-step practical journey, covering everything from the mathematical and theoretical aspects of neural networks, to building your own deep neural networks into your applications with the C# and .NET frameworks. This book begins by giving you a quick refresher of neural networks. You will learn how to build a neural network from scratch using packages such as Encog, Aforge, and Accord. You will learn about various concepts and techniques, such as deep networks, perceptrons, optimization algorithms, convolutional networks, and autoencoders. You will learn ways to add intelligent features to your .NET apps, such as facial and motion detection, object detection and labeling, language understanding, knowledge, and intelligent search. Throughout this book, you will be working on interesting demonstrations that will make it easier to implement complex neural networks in your enterprise applications. What you will learnUnderstand perceptrons and how to implement them in C#Learn how to train and visualize a neural

network using cognitive services. Perform image recognition for detecting and labeling objects using C# and TensorFlowSharp. Detect specific image characteristics such as a face using Accord.Net. Demonstrate particle swarm optimization using a simple XOR problem and EncogTrain convolutional neural networks using ConvNetSharp. Find optimal parameters for your neural network functions using numeric and heuristic optimization techniques.

Who this book is for: This book is for Machine Learning Engineers, Data Scientists, Deep Learning Aspirants and Data Analysts who are now looking to move into advanced machine learning and deep learning with C#. Prior knowledge of machine learning and working experience with C# programming is required to take most out of this book.

### **Deep Learning for Biomedical Image Reconstruction** BALIGE PUBLISHING

**WORKSHOP 1:** In this workshop, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to implement deep learning on recognizing traffic signs using GTSRB dataset, detecting brain tumor using Brain Image MRI dataset, classifying gender, and recognizing facial expression using FER2013 dataset. In Chapter 1, you will learn to create GUI applications to display line graph using PyQt. You will also learn how to display image and its histogram. In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, Pandas, NumPy and other libraries to perform prediction on handwritten digits using MNIST dataset with PyQt. You will build a GUI application for this purpose. In Chapter 3, you will learn how to perform recognizing traffic signs using GTSRB dataset from Kaggle. There are several different types of traffic signs like speed limits, no entry, traffic signals, turn left or right, children crossing, no passing of heavy vehicles, etc. Traffic signs classification is the process of identifying which class a traffic sign belongs to. In this Python project, you will build a deep neural network model that can classify traffic signs in image into different categories. With this model, you will be able to read and understand traffic signs which are a very important task for all autonomous vehicles. You will build a GUI application for this purpose. In Chapter 4, you will learn how to perform detecting brain tumor using Brain Image MRI dataset provided by Kaggle (<https://www.kaggle.com/navoneel/brain-mri-images-for-brain-tumor-detection>) using CNN model. You will build a GUI application

for this purpose. In Chapter 5, you will learn how to perform classifying gender using dataset provided by Kaggle (<https://www.kaggle.com/cashutosh/gender-classification-dataset>) using MobileNetV2 and CNN models. You will build a GUI application for this purpose. In Chapter 6, you will learn how to perform recognizing facial expression using FER2013 dataset provided by Kaggle (<https://www.kaggle.com/nicolejyt/facialexpressionrecognition>) using CNN model. You will also build a GUI application for this purpose.

**WORKSHOP 2:** In this workshop, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to implement deep learning on classifying fruits, classifying cats/dogs, detecting furnitures, and classifying fashion. In Chapter 1, you will learn to create GUI applications to display line graph using PyQt. You will also learn how to display image and its histogram. Then, you will learn how to use OpenCV, NumPy, and other libraries to perform feature extraction with Python GUI (PyQt). The feature detection techniques used in this chapter are Harris Corner Detection, Shi-Tomasi Corner Detector, and Scale-Invariant Feature Transform (SIFT). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying fruits using Fruits 360 dataset provided by Kaggle (<https://www.kaggle.com/moltean/fruits/code>) using Transfer Learning and CNN models. You will build a GUI application for this purpose. In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying cats/dogs using dataset provided by Kaggle (<https://www.kaggle.com/chetankv/dogs-cats-images>) using CNN with Data Generator. You will build a GUI application for this purpose. In Chapter 4, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting furnitures using Furniture Detector dataset provided by Kaggle (<https://www.kaggle.com/akkithetechie/furniture-detector>) using VGG16 model. You will build a GUI application for this purpose. In Chapter 5, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying fashion using Fashion MNIST dataset provided by Kaggle (<https://www.kaggle.com/zalando-research/fashionmnist/code>)

using CNN model. You will build a GUI application for this purpose.

**WORKSHOP 3:** In this workshop, you will implement deep learning on detecting vehicle license plates, recognizing sign language, and detecting surface crack using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting vehicle license plates using Car License Plate Detection dataset provided by Kaggle (<https://www.kaggle.com/andrewmvd/car-plate-detection/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform sign language recognition using Sign Language Digits Dataset provided by Kaggle (<https://www.kaggle.com/ardamavi/sign-language-digits-dataset/download>). In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting surface crack using Surface Crack Detection provided by Kaggle (<https://www.kaggle.com/arunrk7/surface-crack-detection/download>).

**WORKSHOP 4:** In this workshop, implement deep learning-based image classification on detecting face mask, classifying weather, and recognizing flower using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting face mask using Face Mask Detection Dataset provided by Kaggle (<https://www.kaggle.com/omkargurav/face-mask-dataset/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify weather using Multi-class Weather Dataset provided by Kaggle (<https://www.kaggle.com/pratik2901/multiclass-weather-dataset/download>).

**WORKSHOP 5:** In this workshop, implement deep learning-based image classification on classifying monkey species, recognizing rock, paper, and scissor, and classify airplane, car, and ship using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify monkey species using 10 Monkey Species dataset provided by

Kaggle (<https://www.kaggle.com/slothkong/10-monkey-species/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to recognize rock, paper, and scissor using 10 Monkey Species dataset provided by Kaggle (<https://www.kaggle.com/sanikamal/rock-paper-scissors-dataset/download>). WORKSHOP 6: In this worksshop, you will implement two data science projects using Scikit-Learn, Scipy, and other libraries with Python GUI. In Chapter 1, you will learn how to use Scikit-Learn, Scipy, and other libraries to perform how to predict traffic (number of vehicles) in four different junctions using Traffic Prediction Dataset provided by Kaggle (<https://www.kaggle.com/fedesoriano/traffic-prediction-dataset/download>). This dataset contains 48.1k (48120) observations of the number of vehicles each hour in four different junctions: 1) DateTime; 2) Junction; 3) Vehicles; and 4) ID. In Chapter 2, you will learn how to use Scikit-Learn, NumPy, Pandas, and other libraries to perform how to analyze and predict heart attack using Heart Attack Analysis & Prediction Dataset provided by Kaggle (<https://www.kaggle.com/rashikrahmanpritom/heart-attack-analysis-prediction-dataset/download>). WORKSHOP 7: In this workshop, you will implement two data science projects using Scikit-Learn, Scipy, and other libraries with Python GUI. In Project 1, you will learn how to use Scikit-Learn, NumPy, Pandas, Seaborn, and other libraries to perform how to predict early stage diabetes using Early Stage Diabetes Risk Prediction Dataset provided by Kaggle (<https://www.kaggle.com/ishandutta/early-stage-diabetes-risk-prediction-dataset/download>). This dataset contains the sign and symptpom data of newly diabetic or would be diabetic patient. This has been collected using direct questionnaires from the patients of Sylhet Diabetes Hospital in Sylhet, Bangladesh and approved by a doctor. You will develop a GUI using PyQt5 to plot distribution of features, feature importance, cross validation score, and prediced values versus true values. The machine learning models used in this project are Adaboost, Random Forest, Gradient Boosting, Logistic Regression, and Support Vector Machine. In Project 2, you will learn how to use Scikit-Learn, NumPy, Pandas, and other libraries to perform how to analyze and predict breast cancer using Breast Cancer Prediction Dataset provided by Kaggle

(<https://www.kaggle.com/merishnasuwal/breast-cancer-prediction-dataset/download>). Worldwide, breast cancer is the most common type of cancer in women and the second highest in terms of mortality rates. Diagnosis of breast cancer is performed when an abnormal lump is found (from self-examination or x-ray) or a tiny speck of calcium is seen (on an x-ray). After a suspicious lump is found, the doctor will conduct a diagnosis to determine whether it is cancerous and, if so, whether it has spread to other parts of the body. This breast cancer dataset was obtained from the University of Wisconsin Hospitals, Madison from Dr. William H. Wolberg. You will develop a GUI using PyQt5 to plot distribution of features, pairwise relationship, test scores, prediced values versus true values, confusion matrix, and decision boundary. The machine learning models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, and Support Vector Machine. WORKSHOP 8: In this workshop, you will learn how to use Scikit-Learn, TensorFlow, Keras, NumPy, Pandas, Seaborn, and other libraries to implement brain tumor classification and detection with machine learning using Brain Tumor dataset provided by Kaggle. This dataset contains five first order features: Mean (the contribution of individual pixel intensity for the entire image), Variance (used to find how each pixel varies from the neighboring pixel 0, Standard Deviation (the deviation of measured Values or the data from its mean), Skewness (measures of symmetry), and Kurtosis (describes the peak of e.g. a frequency distribution). It also contains eight second order features: Contrast, Energy, ASM (Angular second moment), Entropy, Homogeneity, Dissimilarity, Correlation, and Coarseness. The machine learning models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, and Support Vector Machine. The deep learning models used in this project are MobileNet and ResNet50. In this project, you will develop a GUI using PyQt5 to plot boundary decision, ROC, distribution of features, feature importance, cross validation score, and predicted values versus true values, confusion matrix, training loss, and training accuracy. WORKSHOP 9: In this workshop, you will learn how to use Scikit-Learn, Keras, TensorFlow, NumPy, Pandas, Seaborn, and other libraries to perform COVID-19 Epitope Prediction using COVID-19/SARS B-cell Epitope Prediction dataset provided in Kaggle. All of three datasets consists of information of

protein and peptide: parent\_protein\_id : parent protein ID; protein\_seq : parent protein sequence; start\_position : start position of peptide; end\_position : end position of peptide; peptide\_seq : peptide sequence; chou\_fasman : peptide feature; emini : peptide feature, relative surface accessibility; kolaskar\_tongaonkar : peptide feature, antigenicity; parker : peptide feature, hydrophobicity; isoelectric\_point : protein feature; aromacity: protein feature; hydrophobicity : protein feature; stability : protein feature; and target : antibody valence (target value). The machine learning models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, Support Vector Machine, Adaboost, Gradient Boosting, XGB classifier, and MLP classifier. Then, you will learn how to use sequential CNN and VGG16 models to detect and predict Covid-19 X-RAY using COVID-19 Xray Dataset (Train & Test Sets) provided in Kaggle. The folder itself consists of two subfolders: test and train. Finally, you will develop a GUI using PyQt5 to plot boundary decision, ROC, distribution of features, feature importance, cross validation score, and predicted values versus true values, confusion matrix, training loss, and training accuracy. WORKSHOP 10: In this workshop, you will learn how to use Scikit-Learn, Keras, TensorFlow, NumPy, Pandas, Seaborn, and other libraries to perform analyzing and predicting stroke using dataset provided in Kaggle. The dataset consists of attribute information: id: unique identifier; gender: "Male", "Female" or "Other"; age: age of the patient; hypertension: 0 if the patient doesn't have hypertension, 1 if the patient has hypertension; heart\_disease: 0 if the patient doesn't have any heart diseases, 1 if the patient has a heart disease; ever\_married: "No" or "Yes"; work\_type: "children", "Govt\_jov", "Never\_worked", "Private" or "Self-employed"; Residence\_type: "Rural" or "Urban"; avg\_glucose\_level: average glucose level in blood; bmi: body mass index; smoking\_status: "formerly smoked", "never smoked", "smokes" or "Unknown"; and stroke: 1 if the patient had a stroke or 0 if not. The models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, Support Vector Machine, Adaboost, LGBM classifier, Gradient Boosting, XGB classifier, MLP classifier, and CNN 1D. Finally, you will develop a GUI using PyQt5 to plot boundary decision, ROC, distribution of features, feature importance, cross validation score, and predicted values versus true values,

confusion matrix, learning curve, performance of the model, scalability of the model, training loss, and training accuracy.

**WORKSHOP 11:** In this workshop, you will learn how to use Scikit-Learn, Keras, TensorFlow, NumPy, Pandas, Seaborn, and other libraries to perform classifying and predicting Hepatitis C using dataset provided by UCI Machine Learning Repository. All attributes in dataset except Category and Sex are numerical. Attributes 1 to 4 refer to the data of the patient: X (Patient ID/No.), Category (diagnosis) (values: '0=Blood Donor', '0s=suspect Blood Donor', '1=Hepatitis', '2=Fibrosis', '3=Cirrhosis'), Age (in years), Sex (f,m), ALB, ALP, ALT, AST, BIL, CHE, CHOL, CREA, GGT, and PROT. The target attribute for classification is Category (2): blood donors vs. Hepatitis C patients (including its progress ('just' Hepatitis C, Fibrosis, Cirrhosis)). The models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, Support Vector Machine, Adaboost, LGBM classifier, Gradient Boosting, XGB classifier, MLP classifier, and ANN 1D. Finally, you will develop a GUI using PyQt5 to plot boundary decision, ROC, distribution of features, feature importance, cross validation score, and predicted values versus true values, confusion matrix, learning curve, performance of the model, scalability of the model, training loss, and training accuracy.

[DATA SCIENCE WORKSHOP: Chronic Kidney Disease Classification and Prediction Using Machine Learning and Deep Learning with Python GUI](#) BALIGE PUBLISHING

Big Data Analytics examines large amounts of data to uncover hidden patterns, correlations and other insights. MATLAB has the tool Neural Network Toolbox (Deep Learning Toolbox from version 18) that provides algorithms, functions, and apps to create, train, visualize, and simulate neural networks. You can perform classification, regression, clustering, dimensionality reduction, time-series forecasting, and dynamic system modeling and control. The toolbox includes convolutional neural network and autoencoder deep learning algorithms for image classification and feature learning tasks. To speed up training of large data sets, you can distribute computations and data across multicore processors, GPUs, and computer clusters using Big Data tools (Parallel Computing Toolbox). Unsupervised learning algorithms, including self-organizing maps and competitive layers-Apps for data-fitting, pattern recognition, and clustering-Preprocessing,

postprocessing, and network visualization for improving training efficiency and assessing network performance. his book develops cluster analysis and pattern recognition

**DATA SCIENCE WORKSHOP: Alzheimer's Disease Classification and Prediction Using Machine Learning and Deep Learning with Python GUI** Springer

In the context of sentiment analysis and opinion mining, this project began with dataset exploration. The dataset, comprising user reviews or social media posts, was examined to understand the sentiment labels' distribution. This analysis provided insights into the prevalence of positive or negative opinions, laying the foundation for sentiment classification. To tackle sentiment classification, we employed a range of machine learning algorithms, including Support Vector, Logistic Regression, K-Nearest Neighbours Classifier, Decision Tree, Random Forest Classifier, Gradient Boosting, Extreme Gradient Boosting, Light Gradient Boosting, and Adaboost Classifiers. These algorithms were combined with different vectorization techniques such as Hashing Vectorizer, Count Vectorizer, and TF-IDF Vectorizer. By converting text data into numerical representations, these models were trained and evaluated to identify the most effective combination for sentiment classification. In addition to traditional machine learning algorithms, we explored the power of recurrent neural networks (RNNs) and their variant, Long Short-Term Memory (LSTM). LSTM is particularly adept at capturing contextual dependencies and handling sequential data. The text data was tokenized and padded to ensure consistent input length, allowing the LSTM model to learn from the sequential nature of the text. Performance metrics, including accuracy, were used to evaluate the model's ability to classify sentiments accurately. Furthermore, we delved into Convolutional Neural Networks (CNNs), another deep learning model known for its ability to extract meaningful features. The text data was preprocessed and transformed into numerical representations suitable for CNN input. The architecture of the CNN model, consisting of embedding, convolutional, pooling, and dense layers, facilitated the extraction of relevant features and the classification of sentiments. Analyzing the results of our machine learning models, we gained insights into their effectiveness in sentiment classification. We observed the accuracy and performance of various algorithms and vectorization techniques, enabling us to

identify the models that achieved the highest accuracy and overall performance. LSTM and CNN, being more advanced models, aimed to capture complex patterns and dependencies in the text data, potentially resulting in improved sentiment classification. Monitoring the training history and metrics of the LSTM and CNN models provided valuable insights. We examined the learning progress, convergence behavior, and generalization capabilities of the models. Through the evaluation of performance metrics and convergence trends, we gained an understanding of the models' ability to learn from the data and make accurate predictions. Confusion matrices played a crucial role in assessing the models' predictions. They provided a detailed analysis of the models' classification performance, highlighting the distribution of correct and incorrect classifications for each sentiment category. This analysis allowed us to identify potential areas of improvement and fine-tune the models accordingly. In addition to confusion matrices, visualizations comparing the true values with the predicted values were employed to evaluate the models' performance. These visualizations provided a comprehensive overview of the models' classification accuracy and potential areas for improvement. They allowed us to assess the alignment between the models' predictions and the actual sentiment labels, enabling a deeper understanding of the models' strengths and weaknesses. Overall, the exploration of machine learning, LSTM, and CNN models for sentiment analysis and opinion mining aimed to develop effective tools for understanding public opinions. The results obtained from this project showcased the models' performance, convergence behavior, and their ability to accurately classify sentiments. These insights can be leveraged by businesses and organizations to gain a deeper understanding of the sentiments expressed towards their products or services, enabling them to make informed decisions and adapt their strategies accordingly.

[Hands-On Machine Learning with C#](#) BPB Publications

In this book, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to implement deep learning on recognizing traffic signs using GTSRB dataset, detecting brain tumor using Brain Image MRI dataset, classifying gender, and recognizing facial expression using FER2013 dataset. In Chapter 1, you will learn to create GUI applications to display image histogram. It is a graphical representation that displays the

distribution of pixel intensities in an image. It provides information about the frequency of occurrence of each intensity level in the image. The histogram allows us to understand the overall brightness or contrast of the image and can reveal important characteristics such as dynamic range, exposure, and the presence of certain image features. In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, Pandas, NumPy and other libraries to perform prediction on handwritten digits using MNIST dataset. The MNIST dataset is a widely used dataset in machine learning and computer vision, particularly for image classification tasks. It consists of a collection of handwritten digits from zero to nine, where each digit is represented as a 28x28 grayscale image. The dataset was created by collecting handwriting samples from various individuals and then preprocessing them to standardize the format. Each image in the dataset represents a single digit and is labeled with the corresponding digit it represents. The labels range from 0 to 9, indicating the true value of the handwritten digit. In Chapter 3, you will learn how to perform recognizing traffic signs using GTSRB dataset from Kaggle. There are several different types of traffic signs like speed limits, no entry, traffic signals, turn left or right, children crossing, no passing of heavy vehicles, etc. Traffic signs classification is the process of identifying which class a traffic sign belongs to. In this Python project, you will build a deep neural network model that can classify traffic signs in image into different categories. With this model, you will be able to read and understand traffic signs which are a very important task for all autonomous vehicles. You will build a GUI application for this purpose. In Chapter 4, you will learn how to perform detecting brain tumor using Brain Image MRI dataset. Following are the steps taken in this chapter: Dataset Exploration: Explore the Brain Image MRI dataset from Kaggle. Describe the structure of the dataset, the different classes (tumor vs. non-tumor), and any preprocessing steps required; Data Preprocessing: Preprocess the dataset to prepare it for model training. This may include tasks such as resizing images, normalizing pixel values, splitting data into training and testing sets, and creating labels; Model Building: Use TensorFlow and Keras to build a deep learning model for brain tumor detection. Choose an appropriate architecture, such as a convolutional neural network (CNN), and configure the model layers; Model Training: Train the brain tumor detection model

using the preprocessed dataset. Specify the loss function, optimizer, and evaluation metrics. Monitor the training process and visualize the training/validation accuracy and loss over epochs; Model Evaluation: Evaluate the trained model on the testing dataset. Calculate metrics such as accuracy, precision, recall, and F1 score to assess the model's performance; Prediction and Visualization: Use the trained model to make predictions on new MRI images. Visualize the predicted results alongside the ground truth labels to demonstrate the effectiveness of the model. Finally, you will build a GUI application for this purpose. In Chapter 5, you will learn how to perform classifying gender using dataset provided by Kaggle using MobileNetV2 and CNN models. Following are the steps taken in this chapter: Data Exploration: Load the dataset using Pandas, perform exploratory data analysis (EDA) to gain insights into the data, and visualize the distribution of gender classes; Data Preprocessing: Preprocess the dataset by performing necessary transformations, such as resizing images, converting labels to numerical format, and splitting the data into training, validation, and test sets; Model Building: Use TensorFlow and Keras to build a gender classification model. Define the architecture of the model, compile it with appropriate loss and optimization functions, and summarize the model's structure; Model Training: Train the model on the training set, monitor its performance on the validation set, and tune hyperparameters if necessary. Visualize the training history to analyze the model's learning progress; Model Evaluation: Evaluate the trained model's performance on the test set using various metrics such as accuracy, precision, recall, and F1 score. Generate a classification report and a confusion matrix to assess the model's performance in detail; Prediction and Visualization: Use the trained model to make gender predictions on new, unseen data. Visualize a few sample predictions along with the corresponding images. Finally, you will build a GUI application for this purpose. In Chapter 6, you will learn how to perform recognizing facial expression using FER2013 dataset using CNN model. The FER2013 dataset contains facial images categorized into seven different emotions: anger, disgust, fear, happiness, sadness, surprise, and neutral. To perform facial expression recognition using this dataset, you would typically follow these steps; Data Preprocessing: Load and preprocess the dataset. This may involve resizing the images, converting them to grayscale, and normalizing the pixel values;

Data Split: Split the dataset into training, validation, and testing sets. The training set is used to train the model, the validation set is used to tune hyperparameters and evaluate the model's performance during training, and the testing set is used to assess the final model's accuracy; Model Building: Build a deep learning model using TensorFlow and Keras. This typically involves defining the architecture of the model, selecting appropriate layers (such as convolutional layers, pooling layers, and fully connected layers), and specifying the activation functions and loss functions; Model Training: Train the model using the training set. This involves feeding the training images through the model, calculating the loss, and updating the model's parameters using optimization techniques like backpropagation and gradient descent; Model Evaluation: Evaluate the trained model's performance using the validation set. This can include calculating metrics such as accuracy, precision, recall, and F1 score to assess how well the model is performing; Model Testing: Assess the model's accuracy and performance on the testing set, which contains unseen data. This step helps determine how well the model generalizes to new, unseen facial expressions; Prediction: Use the trained model to make predictions on new images or live video streams. This involves detecting faces in the images using OpenCV, extracting facial features, and feeding the processed images into the model for prediction. Then, you will also build a GUI application for this purpose.

[Deep Belief Nets in C++ and Cuda C](#) BALIGE PUBLISHING BOOK 1: LEARN FROM SCRATCH MACHINE LEARNING WITH PYTHON GUI In this book, you will learn how to use NumPy, Pandas, OpenCV, Scikit-Learn and other libraries to how to plot graph and to process digital image. Then, you will learn how to classify features using Perceptron, Adaline, Logistic Regression (LR), Support Vector Machine (SVM), Decision Tree (DT), Random Forest (RF), and K-Nearest Neighbor (KNN) models. You will also learn how to extract features using Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), Kernel Principal Component Analysis (KPCA) algorithms and use them in machine learning. In Chapter 1, you will learn: Tutorial Steps To Create A Simple GUI Application, Tutorial Steps to Use Radio Button, Tutorial Steps to Group Radio Buttons, Tutorial Steps to Use CheckBox Widget, Tutorial Steps to Use Two CheckBox Groups, Tutorial Steps to Understand Signals and Slots, Tutorial Steps to

Convert Data Types, Tutorial Steps to Use Spin Box Widget, Tutorial Steps to Use ScrollBar and Slider, Tutorial Steps to Use List Widget, Tutorial Steps to Select Multiple List Items in One List Widget and Display It in Another List Widget, Tutorial Steps to Insert Item into List Widget, Tutorial Steps to Use Operations on Widget List, Tutorial Steps to Use Combo Box, Tutorial Steps to Use Calendar Widget and Date Edit, and Tutorial Steps to Use Table Widget. In Chapter 2, you will learn: Tutorial Steps To Create A Simple Line Graph, Tutorial Steps To Create A Simple Line Graph in Python GUI, Tutorial Steps To Create Two or More Graphs in the Same Axis, Tutorial Steps To Create Two Axes in One Canvas, Tutorial Steps To Use Two Widgets, Tutorial Steps To Use Two Widgets, Each of Which Has Two Axes, Tutorial Steps To Use Axes With Certain Opacity Levels, Tutorial Steps To Choose Line Color From Combo Box, Tutorial Steps To Calculate Fast Fourier Transform, Tutorial Steps To Create GUI For FFT, Tutorial Steps To Create GUI For FFT With Some Other Input Signals, Tutorial Steps To Create GUI For Noisy Signal, Tutorial Steps To Create GUI For Noisy Signal Filtering, and Tutorial Steps To Create GUI For Wav Signal Filtering. In Chapter 3, you will learn: Tutorial Steps To Convert RGB Image Into Grayscale, Tutorial Steps To Convert RGB Image Into YUV Image, Tutorial Steps To Convert RGB Image Into HSV Image, Tutorial Steps To Filter Image, Tutorial Steps To Display Image Histogram, Tutorial Steps To Display Filtered Image Histogram, Tutorial Steps To Filter Image With CheckBoxes, Tutorial Steps To Implement Image Thresholding, and Tutorial Steps To Implement Adaptive Image Thresholding. You will also learn: Tutorial Steps To Generate And Display Noisy Image, Tutorial Steps To Implement Edge Detection On Image, Tutorial Steps To Implement Image Segmentation Using Multiple Thresholding and K-Means Algorithm, Tutorial Steps To Implement Image Denoising, Tutorial Steps To Detect Face, Eye, and Mouth Using Haar Cascades, Tutorial Steps To Detect Face Using Haar Cascades with PyQt, Tutorial Steps To Detect Eye, and Mouth Using Haar Cascades with PyQt, Tutorial Steps To Extract Detected Objects, Tutorial Steps To Detect Image Features Using Harris Corner Detection, Tutorial Steps To Detect Image Features Using Shi-Tomasi Corner Detection, Tutorial Steps To Detect Features Using Scale-Invariant Feature Transform (SIFT), and Tutorial Steps To Detect Features Using

Features from Accelerated Segment Test (FAST). In Chapter 4, In this tutorial, you will learn how to use Pandas, NumPy and other libraries to perform simple classification using perceptron and Adaline (adaptive linear neuron). The dataset used is Iris dataset directly from the UCI Machine Learning Repository. You will learn: Tutorial Steps To Implement Perceptron, Tutorial Steps To Implement Perceptron with PyQt, Tutorial Steps To Implement Adaline (ADaptive LInear NEuron), and Tutorial Steps To Implement Adaline with PyQt. In Chapter 5, you will learn how to use the scikit-learn machine learning library, which provides a wide variety of machine learning algorithms via a user-friendly Python API and to perform classification using perceptron, Adaline (adaptive linear neuron), and other models. The dataset used is Iris dataset directly from the UCI Machine Learning Repository. You will learn: Tutorial Steps To Implement Perceptron Using Scikit-Learn, Tutorial Steps To Implement Perceptron Using Scikit-Learn with PyQt, Tutorial Steps To Implement Logistic Regression Model, Tutorial Steps To Implement Logistic Regression Model with PyQt, Tutorial Steps To Implement Logistic Regression Model Using Scikit-Learn with PyQt, Tutorial Steps To Implement Support Vector Machine (SVM) Using Scikit-Learn, Tutorial Steps To Implement Decision Tree (DT) Using Scikit-Learn, Tutorial Steps To Implement Random Forest (RF) Using Scikit-Learn, and Tutorial Steps To Implement K-Nearest Neighbor (KNN) Using Scikit-Learn. In Chapter 6, you will learn how to use Pandas, NumPy, Scikit-Learn, and other libraries to implement different approaches for reducing the dimensionality of a dataset using different feature selection techniques. You will learn about three fundamental techniques that will help us to summarize the information content of a dataset by transforming it onto a new feature subspace of lower dimensionality than the original one. Data compression is an important topic in machine learning, and it helps us to store and analyze the increasing amounts of data that are produced and collected in the modern age of technology. You will learn the following topics: Principal Component Analysis (PCA) for unsupervised data compression, Linear Discriminant Analysis (LDA) as a supervised dimensionality reduction technique for maximizing class separability, Nonlinear dimensionality reduction via Kernel Principal Component Analysis (KPCA). You will learn: Tutorial Steps To Implement Principal Component Analysis (PCA), Tutorial Steps To Implement Principal Component Analysis (PCA)

Using Scikit-Learn, Tutorial Steps To Implement Principal Component Analysis (PCA) Using Scikit-Learn with PyQt, Tutorial Steps To Implement Linear Discriminant Analysis (LDA), Tutorial Steps To Implement Linear Discriminant Analysis (LDA) with Scikit-Learn, Tutorial Steps To Implement Linear Discriminant Analysis (LDA) Using Scikit-Learn with PyQt, Tutorial Steps To Implement Kernel Principal Component Analysis (KPCA) Using Scikit-Learn, and Tutorial Steps To Implement Kernel Principal Component Analysis (KPCA) Using Scikit-Learn with PyQt. In Chapter 7, you will learn how to use Keras, Scikit-Learn, Pandas, NumPy and other libraries to perform prediction on handwritten digits using MNIST dataset. You will learn: Tutorial Steps To Load MNIST Dataset, Tutorial Steps To Load MNIST Dataset with PyQt, Tutorial Steps To Implement Perceptron With PCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Perceptron With LDA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Perceptron With KPCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Logistic Regression (LR) Model With PCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Logistic Regression (LR) Model With LDA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Logistic Regression (LR) Model With KPCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement , Tutorial Steps To Implement Support Vector Machine (SVM) Model With LDA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Support Vector Machine (SVM) Model With KPCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Decision Tree (DT) Model With PCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Decision Tree (DT) Model With LDA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Decision Tree (DT) Model With KPCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Random Forest (RF) Model With PCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Random Forest (RF) Model With LDA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement Random Forest (RF) Model With KPCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To Implement K-Nearest Neighbor (KNN) Model With PCA Feature Extractor on MNIST Dataset Using PyQt, Tutorial Steps To

Implement K-Nearest Neighbor (KNN) Model With LDA Feature Extractor on MNIST Dataset Using PyQt, and Tutorial Steps To Implement K-Nearest Neighbor (KNN) Model With KPCA Feature Extractor on MNIST Dataset Using PyQt. BOOK 2: THE PRACTICAL GUIDES ON DEEP LEARNING USING SCIKIT-LEARN, KERAS, AND TENSORFLOW WITH PYTHON GUI In this book, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to implement deep learning on recognizing traffic signs using GTSRB dataset, detecting brain tumor using Brain Image MRI dataset, classifying gender, and recognizing facial expression using FER2013 dataset In Chapter 1, you will learn to create GUI applications to display line graph using PyQt. You will also learn how to display image and its histogram. In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, Pandas, NumPy and other libraries to perform prediction on handwritten digits using MNIST dataset with PyQt. You will build a GUI application for this purpose. In Chapter 3, you will learn how to perform recognizing traffic signs using GTSRB dataset from Kaggle. There are several different types of traffic signs like speed limits, no entry, traffic signals, turn left or right, children crossing, no passing of heavy vehicles, etc. Traffic signs classification is the process of identifying which class a traffic sign belongs to. In this Python project, you will build a deep neural network model that can classify traffic signs in image into different categories. With this model, you will be able to read and understand traffic signs which are a very important task for all autonomous vehicles. You will build a GUI application for this purpose. In Chapter 4, you will learn how to perform detecting brain tumor using Brain Image MRI dataset provided by Kaggle (<https://www.kaggle.com/navoneel/brain-mri-images-for-brain-tumor-detection>) using CNN model. You will build a GUI application for this purpose. In Chapter 5, you will learn how to perform classifying gender using dataset provided by Kaggle (<https://www.kaggle.com/cashutosh/gender-classification-dataset>) using MobileNetV2 and CNN models. You will build a GUI application for this purpose. In Chapter 6, you will learn how to perform recognizing facial expression using FER2013 dataset provided by Kaggle (<https://www.kaggle.com/nicolejyt/facialexpressionrecognition>) using CNN model. You will also build a GUI application for this purpose. BOOK 3: STEP BY STEP TUTORIALS ON DEEP LEARNING

USING SCIKIT-LEARN, KERAS, AND TENSORFLOW WITH PYTHON GUI In this book, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to implement deep learning on classifying fruits, classifying cats/dogs, detecting furnitures, and classifying fashion. In Chapter 1, you will learn to create GUI applications to display line graph using PyQt. You will also learn how to display image and its histogram. Then, you will learn how to use OpenCV, NumPy, and other libraries to perform feature extraction with Python GUI (PyQt). The feature detection techniques used in this chapter are Harris Corner Detection, Shi-Tomasi Corner Detector, and Scale-Invariant Feature Transform (SIFT). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying fruits using Fruits 360 dataset provided by Kaggle (<https://www.kaggle.com/moltean/fruits/code>) using Transfer Learning and CNN models. You will build a GUI application for this purpose. In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying cats/dogs using dataset provided by Kaggle (<https://www.kaggle.com/chetankv/dogs-cats-images>) using CNN with Data Generator. You will build a GUI application for this purpose. In Chapter 4, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting furnitures using Furniture Detector dataset provided by Kaggle (<https://www.kaggle.com/akkithetechie/furniture-detector>) using VGG16 model. You will build a GUI application for this purpose. In Chapter 5, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform classifying fashion using Fashion MNIST dataset provided by Kaggle (<https://www.kaggle.com/zalando-research/fashionmnist/code>) using CNN model. You will build a GUI application for this purpose. BOOK 4: Project-Based Approach On DEEP LEARNING Using Scikit-Learn, Keras, And TensorFlow with Python GUI In this book, implement deep learning on detecting vehicle license plates, recognizing sign language, and detecting surface crack using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to

perform detecting vehicle license plates using Car License Plate Detection dataset provided by Kaggle (<https://www.kaggle.com/andrewmvd/car-plate-detection/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform sign language recognition using Sign Language Digits Dataset provided by Kaggle (<https://www.kaggle.com/ardamavi/sign-language-digits-dataset/download>). In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting surface crack using Surface Crack Detection provided by Kaggle (<https://www.kaggle.com/arunrk7/surface-crack-detection/download>). BOOK 5: Hands-On Guide To IMAGE CLASSIFICATION Using Scikit-Learn, Keras, And TensorFlow with PYTHON GUI In this book, implement deep learning-based image classification on detecting face mask, classifying weather, and recognizing flower using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting face mask using Face Mask Detection Dataset provided by Kaggle (<https://www.kaggle.com/omkargurav/face-mask-dataset/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify weather using Multi-class Weather Dataset provided by Kaggle (<https://www.kaggle.com/pratik2901/multiclass-weather-dataset/download>). In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to recognize flower using Flowers Recognition dataset provided by Kaggle (<https://www.kaggle.com/alxmamaev/flowers-recognition/download>). BOOK 6: Step by Step Tutorial IMAGE CLASSIFICATION Using Scikit-Learn, Keras, And TensorFlow with PYTHON GUI In this book, implement deep learning-based image classification on classifying monkey species, recognizing rock, paper, and scissor, and classify airplane, car, and ship using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In Chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify

monkey species using 10 Monkey Species dataset provided by Kaggle (<https://www.kaggle.com/slothkong/10-monkey-species/download>). In Chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to recognize rock, paper, and scissor using 10 Monkey Species dataset provided by Kaggle (<https://www.kaggle.com/sanikamal/rock-paper-scissors-dataset/download>). In Chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify airplane, car, and ship using Multiclass-image-dataset-airplane-car-ship dataset provided by Kaggle (<https://www.kaggle.com/abtabm/multiclassimagedatasetairplane-car>).

*Step by Step Tutorial IMAGE CLASSIFICATION Using Scikit-Learn, Keras, And TensorFlow with PYTHON GUI* Packt Publishing

Apply deep learning techniques and neural network methodologies to build, train, and optimize generative network models

**Key Features**

- Implement GAN architectures to generate images, text, audio, 3D models, and more
- Understand how GANs work and become an active contributor in the open source community
- Learn how to generate photo-realistic images based on text descriptions

**Book Description** With continuously evolving research and development, Generative Adversarial Networks (GANs) are the next big thing in the field of deep learning. This book highlights the key improvements in GANs over generative models and guides in making the best out of GANs with the help of hands-on examples. This book starts by taking you through the core concepts necessary to understand how each component of a GAN model works. You'll build your first GAN model to understand how generator and discriminator networks function. As you advance, you'll delve into a range of examples and datasets to build a variety of GAN networks using PyTorch functionalities and services, and become well-versed with architectures, training strategies, and evaluation methods for image generation, translation, and restoration. You'll even learn how to apply GAN models to solve problems in areas such as computer vision, multimedia, 3D models, and natural language processing (NLP). The book covers how to overcome the challenges faced while building generative models from scratch. Finally, you'll also discover how to train your GAN models to generate adversarial

examples to attack other CNN and GAN models. By the end of this book, you will have learned how to build, train, and optimize next-generation GAN models and use them to solve a variety of real-world problems. What you will learn

- Implement PyTorch's latest features to ensure efficient model designing
- Get to grips with the working mechanisms of GAN models
- Perform style transfer between unpaired image collections with CycleGAN
- Build and train 3D-GANs to generate a point cloud of 3D objects
- Create a range of GAN models to perform various image synthesis operations
- Use SEGAN to suppress noise and improve the quality of speech audio

Who this book is for This GAN book is for machine learning practitioners and deep learning researchers looking to get hands-on guidance in implementing GAN models using PyTorch. You'll become familiar with state-of-the-art GAN architectures with the help of real-world examples. Working knowledge of Python programming language is necessary to grasp the concepts covered in this book.

#### **Neural Networks and Deep Learning** Apress

In this book, implement deep learning-based image classification on classifying monkey species, recognizing rock, paper, and scissor, and classify airplane, car, and ship using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify monkey species using 10 Monkey Species dataset provided by Kaggle (<https://www.kaggle.com/slothkong/10-monkey-species/download>). Here's an overview of the steps involved in classifying monkey species using the 10 Monkey Species dataset:

- Dataset Preparation:** Download the 10 Monkey Species dataset from Kaggle and extract the files. The dataset should consist of separate folders for each monkey species, with corresponding images.
- Load and Preprocess Images:** Use libraries such as OpenCV to load the images from the dataset. Resize the images to a consistent size (e.g., 224x224 pixels) to ensure uniformity.
- Split the Dataset:** Divide the dataset into training and testing sets. Typically, an 80:20 or 70:30 split is used, where the larger portion is used for training and the smaller portion for testing the model's performance.
- Label Encoding:** Encode the categorical labels (monkey species) into numeric form. This step is necessary to train a machine learning model, as most algorithms expect

- numerical inputs.
- Feature Extraction:** Extract meaningful features from the images using techniques like deep learning or image processing algorithms. This step helps in representing the images in a format that the machine learning model can understand.
- Model Training:** Use libraries like TensorFlow and Keras to train a machine learning model on the preprocessed data. Choose an appropriate model architecture, in this case, MobileNetV2.
- Model Evaluation:** Evaluate the trained model on the testing set to assess its performance. Metrics like accuracy, precision, recall, and F1-score can be used to evaluate the model's classification performance.
- Predictions:** Use the trained model to make predictions on new, unseen images. Pass the images through the trained model and obtain the predicted labels for the monkey species. In chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to recognize rock, paper, and scissor using dataset provided by Kaggle (<https://www.kaggle.com/sanikamal/rock-paper-scissors-dataset/download>). Here's the outline of the steps:

- Step 1: Dataset Preparation:** Download the rock-paper-scissors dataset from Kaggle by visiting the provided link and clicking on the "Download" button. Save the dataset to a local directory on your machine. Extract the downloaded dataset to a suitable location. This will create a folder containing the images for rock, paper, and scissors.
- Step 2: Data Preprocessing:** Import the required libraries: TensorFlow, Keras, NumPy, OpenCV, and Pandas. Load the dataset using OpenCV: Iterate through the image files in the dataset directory and use OpenCV's `cv2.imread()` function to load each image. You can specify the image's file extension (e.g., PNG) and directory path. Preprocess the images: Resize the loaded images to a consistent size using OpenCV's `cv2.resize()` function. You may choose a specific width and height suitable for your model. Prepare the labels: Create a list or array to store the corresponding labels for each image (rock, paper, or scissors). This can be done based on the file naming convention or by mapping images to their respective labels using a dictionary.
- Step 3: Model Training:** Create a convolutional neural network (CNN) model using Keras: Define a CNN architecture using Keras' Sequential model or functional API. This typically consists of convolutional layers, pooling layers, and dense layers. Compile the model: Specify the loss function (e.g., categorical cross-



entropy) and optimizer (e.g., Adam) using Keras' compile() function. You can also define additional metrics to evaluate the model's performance. Train the model: Use Keras' fit() function to train the model on the preprocessed dataset. Specify the training data, labels, batch size, number of epochs, and validation data if available. This will optimize the model's weights based on the provided dataset. Save the trained model: Once the model training is complete, you can save the trained model to disk using Keras' save() or save\_weights() function. This allows you to load the model later for predictions or further training.; Step 4: Model Evaluation: Evaluate the trained model: Use Keras' evaluate() function to assess the model's performance on a separate testing dataset. Provide the testing data and labels to calculate metrics such as accuracy, precision, recall, and F1 score. This will help you understand how well the model generalizes to new, unseen data. Analyze the model's performance: Interpret the evaluation metrics and analyze any potential areas of improvement. You can also visualize the confusion matrix or classification report to gain more insights into the model's predictions.; Step 5: Prediction: Use the trained model for predictions: Load the saved model using Keras' load\_model() function. Then, pass new, unseen images through the model to obtain predictions. Preprocess these images in the same way as the training images (resize, normalize, etc.). Visualize and interpret predictions: Display the predicted labels alongside the corresponding images to see how well the model performs. You can use libraries like Matplotlib or OpenCV to show the images and their predicted labels. Additionally, you can calculate the accuracy of the model's predictions on the new dataset. In chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform how to classify airplane, car, and ship using Multiclass-image-dataset-airplane-car-ship dataset provided by Kaggle (<https://www.kaggle.com/abtabm/multiclassimagedatasetairplane-car>). Here are the outline steps: Import the required libraries: TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy. Load and preprocess the dataset: Read the images from the dataset folder. Resize the images to a fixed size. Store the images and corresponding labels.; Split the dataset into training and testing sets: Split the data and labels into training and testing sets using a specified ratio.; Encode the labels: Convert the categorical labels into numerical format. Perform one-hot encoding on the

labels.; Build MobileNetV2 model using Keras: Create a sequential model. Add convolutional layers with activation functions. Add pooling layers for downsampling. Flatten the output and add dense layers. Set the output layer with softmax activation.; Compile and train the model: Compile the model with an optimizer and loss function. Train the model using the training data and labels. Specify the number of epochs and batch size.; Evaluate the model: Evaluate the trained model using the testing data and labels. Calculate the accuracy of the model.; Make predictions on new images: Load and preprocess a new image. Use the trained model to predict the label of the new image. Convert the predicted label from numerical format to categorical.

*OPINION MINING AND PREDICTION USING MACHINE LEARNING AND DEEP LEARNING WITH PYTHON GUI* CESAR PEREZ  
Discover the essential building blocks of the most common forms of deep belief networks. At each step this book provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. The first of three in a series on C++ and CUDA C deep learning and belief nets, *Deep Belief Nets in C++ and CUDA C: Volume 1* shows you how the structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a thought process that is capable of learning abstract concepts built from simpler primitives. As such, you'll see that a typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. All the routines and algorithms presented in the book are available in the code download, which also contains some libraries of related routines. *What You Will Learn* Employ deep learning using C++ and CUDA C Work with supervised feedforward networks Implement restricted Boltzmann machines Use generative samplings Discover why these are important *Who This Book Is For* Those who have at least a basic knowledge of neural networks and some prior programming experience, although some C++ and CUDA C is recommended. *Deep Learning for Coders with fastai and PyTorch* "O'Reilly Media, Inc."

*Big Data Analytics* examines large amounts of data to uncover hidden patterns, correlations and other insights. With today's

technology, it's possible to analyze your data and get answers from it almost immediately - an effort that's slower and less efficient with more traditional business intelligence solutions. Deep learning (also known as deep structured learning, hierarchical learning or deep machine learning) is a branch of machine learning based on a set of algorithms that attempt to model high level abstractions in data. Various deep learning architectures such as deep neural networks, convolutional deep neural networks, deep belief networks and recurrent neural networks have been applied to fields like computer vision, automatic speech recognition, natural language processing, audio recognition and bioinformatics where they have been shown to produce state-of-the-art results on various tasks. Deep learning has been characterized as a buzzword, or a rebranding of neural networks. This book deeps in big data and deep learning techniques

*Computer Vision -- ECCV 2014* BALIGE PUBLISHING

Discover the power of deep neural networks for image reconstruction with this state-of-the-art review of modern theories and applications. The background theory of deep learning is introduced step-by-step, and by incorporating modeling fundamentals this book explains how to implement deep learning in a variety of modalities, including X-ray, CT, MRI and others. Real-world examples demonstrate an interdisciplinary approach to medical image reconstruction processes, featuring numerous imaging applications. Recent clinical studies and innovative research activity in generative models and mathematical theory will inspire the reader towards new frontiers. This book is ideal for graduate students in Electrical or Biomedical Engineering or Medical Physics.

*STOCK PRICE ANALYSIS, PREDICTION, AND FORECASTING USING MACHINE LEARNING AND DEEP LEARNING WITH PYTHON* Packt Publishing Ltd

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, *Deep Learning* is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of

concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

*Deep Belief Nets in C++ and CUDA C: Volume 2* BALIGE PUBLISHING

This dataset is a playground for fundamental and technical analysis. It is said that 30% of traffic on stocks is already generated by machines, can trading be fully automated? If not, there is still a lot to learn from historical data. The dataset consists of data spans from 2010 to the end 2016, for companies new on stock market date range is shorter. To perform forecasting based on regression adjusted closing price of gold, you will use: Linear Regression, Random Forest regression, Decision Tree regression, Support Vector Machine regression, Naïve Bayes regression, K-Nearest Neighbor regression, Adaboost regression, Gradient Boosting regression, Extreme Gradient Boosting regression, Light Gradient Boosting regression, Catboost regression, MLP regression, and LSTM (Long-Short Term Memory)

regression. The machine learning models used predict gold daily returns as target variable are K-Nearest Neighbor classifier, Random Forest classifier, Naive Bayes classifier, Logistic Regression classifier, Decision Tree classifier, Support Vector Machine classifier, LGBM classifier, Gradient Boosting classifier, XGB classifier, MLP classifier, Gaussian Mixture Model classifier, and Extra Trees classifier. Finally, you will plot boundary decision, distribution of features, feature importance, predicted values versus true values, confusion matrix, learning curve, performance of the model, and scalability of the model.

Mastering .NET Machine Learning MIT Press

In this book, implement deep learning on detecting vehicle license plates, recognizing sign language, and detecting surface crack using TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries. In chapter 1, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting vehicle license plates using Car License Plate Detection dataset provided by Kaggle (<https://www.kaggle.com/andrewmvd/car-plate-detection/download>). To perform license plate detection, these steps are taken: 1. Dataset Preparation: Extract the dataset and organize it into separate folders for images and annotations. The annotations should contain bounding box coordinates for license plate regions.; 2. Data Preprocessing: Load the images and annotations from the dataset. Preprocess the images by resizing, normalizing, or applying any other necessary transformations. Convert the annotation bounding box coordinates to the appropriate format for training.; 3. Training Data Generation: Divide the dataset into training and validation sets. Generate training data by augmenting the images and annotations (e.g., flipping, rotating, zooming). Create data generators or data loaders to efficiently load the training data.; 4. Model Development: Choose a suitable deep learning model architecture for license plate detection, such as a convolutional neural network (CNN). Use TensorFlow and Keras to develop the model architecture. Compile the model with appropriate loss functions and optimization algorithms.; 5. Model Training: Train the model using the prepared training data. Monitor the training process by tracking metrics like loss and accuracy. Adjust the hyperparameters or model architecture as needed to improve performance.; 6. Model Evaluation: Evaluate the trained model using the validation set. Calculate relevant

metrics like precision, recall, and F1 score. Make any necessary adjustments to the model based on the evaluation results.; 7. License Plate Detection: Use the trained model to detect license plates in new images. Apply any post-processing techniques to refine the detected regions. Extract the license plate regions and further process them if needed. In chapter 2, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform sign language recognition using Sign Language Digits Dataset. Here are the steps to perform sign language recognition using the Sign Language Digits Dataset: 1. Download the dataset from Kaggle: You can visit the Kaggle Sign Language Digits Dataset page (<https://www.kaggle.com/ardamavi/sign-language-digits-dataset>) and download the dataset.; 2. Extract the dataset: After downloading the dataset, extract the contents from the downloaded zip file to a suitable location on your local machine.; 3. Load the dataset: The dataset consists of two parts - images and a CSV file containing the corresponding labels. The images are stored in a folder, and the CSV file contains the image paths and labels.; 4. Preprocess the dataset: Depending on the specific requirements of your model, you may need to preprocess the dataset. This can include tasks such as resizing images, converting labels to numerical format, normalizing pixel values, or splitting the dataset into training and testing sets.; 5. Build a machine learning model: Use libraries such as TensorFlow and Keras to build a sign language recognition model. This typically involves designing the architecture of the model, compiling it with suitable loss functions and optimizers, and training the model on the preprocessed dataset.; 6. Evaluate the model: After training the model, evaluate its performance using appropriate evaluation metrics. This can help you understand how well the model is performing on the sign language recognition task.; 7. Make predictions: Once the model is trained and evaluated, you can use it to make predictions on new sign language images. Pass the image through the model, and it will predict the corresponding sign language digit. In chapter 3, you will learn how to use TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, NumPy and other libraries to perform detecting surface crack using Surface Crack Detection provided by Kaggle (<https://www.kaggle.com/arunrk7/surface-crack-detection/download>). Here's a general outline of the process: Data Preparation:

Start by downloading the dataset from the Kaggle link you provided. Extract the dataset and organize it into appropriate folders (e.g., training and testing folders).; **Import Libraries:** Begin by importing the necessary libraries, including TensorFlow, Keras, Scikit-Learn, OpenCV, Pandas, and NumPy.; **Data Loading and Preprocessing:** Load the images and labels from the dataset. Since the dataset may come in different formats, it's essential to understand its structure and adjust the code accordingly. Use OpenCV to read the images and Pandas to load the labels.; **Data Augmentation:** Perform data augmentation techniques such as rotation, flipping, and scaling to increase the diversity of the training data and prevent overfitting. You can use the `ImageDataGenerator` class from Keras for this purpose.; **Model Building:** Define your neural network architecture using the Keras API with TensorFlow backend. You can start with a simple architecture like a convolutional neural network (CNN). Experiment with different architectures to achieve better performance.; **Model Compilation:** Compile your model by specifying the loss function, optimizer, and evaluation metric. For a binary classification problem like crack detection, you can use binary cross-entropy as the loss function and Adam as the optimizer.; **Model Training:** Train your model on the prepared dataset using the `fit()` method. Split your data into training and validation sets using `train_test_split()` from Scikit-Learn. Monitor the training progress and adjust hyperparameters as needed. **Model Evaluation:** Evaluate the performance of your trained model on the test set. Use appropriate evaluation metrics such as accuracy, precision, recall, and F1 score. Scikit-Learn provides functions for calculating these metrics.; **Model Prediction:** Use the trained model to predict crack detection on new unseen images. Load the test images, preprocess them if necessary, and use the trained model to make predictions.

**Hands-On Neural Network Programming with C#** Packt Publishing Ltd

This three-volume set LNCS 12888, 12898, and 12890 constitutes the refereed conference proceedings of the 11th International Conference on Image and Graphics, ICIG 2021, held in Haikou, China, in August 2021.\* The 198 full papers presented were selected from 421 submissions and focus on advances of theory, techniques and algorithms as well as innovative technologies of image, video and graphics processing and fostering innovation,

entrepreneurship, and networking. \*The conference was postponed due to the COVID-19 pandemic.

**Image and Graphics** BALIGE PUBLISHING

Gain insights from your data quicker and easier than ever before using your own C#!**About This Book\*** This easy-to-follow guide will help you enter the field of Data Science with C#!**Learn to use various .NET APIs and perform exploratory data analysis that starts with a few data journalism questions and ends with intuitive visualizations\*** Explore the power of C#'s cool Data Science libraries and learn to perform complex computations using this hands-on guide!**Who This Book Is For**If you are a seasoned C# developer with knowledge of the Microsoft technology stack and a good background in mathematics, then this book is for you. A basic understanding of statistics and probability is necessary to get the most out of this book. This book will also be useful for those data scientists who want to know how to perform data analysis with C#.What you will learn\* Be reintroduced to the world of .NET statistics and calculus libraries and frameworks\* Overcome the pitfalls of modeling your data for simple to complex analysis\* Use the tools to utilize the statistical models for designing custom Machine Learning applications\* Know how to deal with high volumes of data and scale your ASP.NET application to match the performance requirements of your Data Science application\* Integrate various tools such as ATOM (ASP.NET MVC), Data Science libraries such as Accord.NET, Python's Natural Language Toolkit, and SWIRL to achieve precise data manipulations and intuitive visualizations\* Implement models such as k-nearest Neighbors, Naive Bayes, linear and logistic regression, decision trees, neural networks, and deep learning**In Detail**Data Science is a new field that is still being defined. It's existed largely in academia and "start-ups," yet stands at the core of the world's most important, fun, and useful software. Because it was traditionally done in an elite academic setting, proprietary software was not included. With Microsoft taking a much more different approach with licensing, Data Science can easily be done with Microsoft tools and is now readily available to C# developers.This book will re-introduce you to the important parts of statistics and probabilities. From there, we move on to how to apply these principles using C# and C# based libraries. This includes using purely Microsoft technologies such as SSRS, SSIS, PowerShell, SSAS, as well as integrating various .NET

libraries available for Data Science.Once you have realized the potential of various libraries, we'll dive deeper in to implementing concepts of regression, inference, correlation, and causation to devise empirical evidence from your data sets and visualize it. Towards the end of the book, you will learn to implement supervised, unsupervised, and deep learning algorithms and techniques that can be applied to a wide range of real-world problems.

**Mobile Artificial Intelligence Projects** BALIGE PUBLISHING

**Book 1: Practical Data Science Programming for Medical Datasets Analysis and Prediction with Python GUI** In this book, you will implement two data science projects using Scikit-Learn, Scipy, and other libraries with Python GUI. In Project 1, you will learn how to use Scikit-Learn, NumPy, Pandas, Seaborn, and other libraries to perform how to predict early stage diabetes using Early Stage Diabetes Risk Prediction Dataset provided by Kaggle. This dataset contains the sign and symptom data of newly diabetic or would be diabetic patient. This has been collected using direct questionnaires from the patients of Sylhet Diabetes Hospital in Sylhet, Bangladesh and approved by a doctor. You will develop a GUI using PyQt5 to plot distribution of features, feature importance, cross validation score, and predicted values versus true values. The machine learning models used in this project are Adaboost, Random Forest, Gradient Boosting, Logistic Regression, and Support Vector Machine. In Project 2, you will learn how to use Scikit-Learn, NumPy, Pandas, and other libraries to perform how to analyze and predict breast cancer using Breast Cancer Prediction Dataset provided by Kaggle. Worldwide, breast cancer is the most common type of cancer in women and the second highest in terms of mortality rates.Diagnosis of breast cancer is performed when an abnormal lump is found (from self-examination or x-ray) or a tiny speck of calcium is seen (on an x-ray). After a suspicious lump is found, the doctor will conduct a diagnosis to determine whether it is cancerous and, if so, whether it has spread to other parts of the body. This breast cancer dataset was obtained from the University of Wisconsin Hospitals, Madison from Dr. William H. Wolberg. You will develop a GUI using PyQt5 to plot distribution of features, pairwise relationship, test scores, predicted values versus true values, confusion matrix, and decision boundary. The machine learning models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes,

Logistic Regression, Decision Tree, and Support Vector Machine. Book 2: Step by Step Tutorials For Data Science With Python GUI: Traffic And Heart Attack Analysis And Prediction In this book, you will implement two data science projects using Scikit-Learn, Scipy, and other libraries with Python GUI. In Chapter 1, you will learn how to use Scikit-Learn, Scipy, and other libraries to perform how to predict traffic (number of vehicles) in four different junctions using Traffic Prediction Dataset provided by Kaggle. This dataset contains 48.1k (48120) observations of the number of vehicles each hour in four different junctions: 1) DateTime; 2) Junction; 3) Vehicles; and 4) ID. In Chapter 2, you will learn how to use Scikit-Learn, NumPy, Pandas, and other libraries to perform how to analyze and predict heart attack using Heart Attack Analysis & Prediction Dataset provided by Kaggle. Book 3: BRAIN TUMOR: Analysis, Classification, and Detection Using Machine Learning and Deep Learning with Python GUI In this project, you will learn how to use Scikit-Learn, TensorFlow, Keras, NumPy, Pandas, Seaborn, and other libraries to implement brain tumor classification and detection with machine learning using Brain Tumor dataset provided by Kaggle. This dataset contains five first order features: Mean (the contribution of individual pixel intensity for the entire image), Variance (used to find how each pixel varies from the neighboring pixel), Standard Deviation (the deviation of measured Values or the data from its mean), Skewness (measures of symmetry), and Kurtosis (describes the peak of e.g. a frequency distribution). It also contains eight second order features: Contrast, Energy, ASM (Angular second moment), Entropy, Homogeneity, Dissimilarity, Correlation, and Coarseness. The machine learning models used in this project are K-Nearest Neighbor, Random Forest, Naive Bayes, Logistic Regression, Decision Tree, and Support Vector Machine. The deep learning models used in this project are MobileNet and ResNet50. In this project, you will develop a GUI using PyQt5 to plot boundary decision, ROC, distribution of features, feature importance, cross validation score, and predicted values versus true values, confusion matrix, training loss, and training accuracy.

**Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow** Packt Publishing Ltd

Get hands on with Kelp.Net, Microsoft's latest Deep Learning framework Key features Deep Learning Basics The ultimate

Kelp.Net reference guide Develop state of the art deep learning applications C# deep learning code Develop advanced deep learning models with minimal code Develop your own advanced deep learning models Loading and Saving Deep Learning Models Comprehensive Kelp.Net reference Sample Deep Learning Models and Tests penCL Reference Easily add deep learning to your applications Many sample models and tests Intuitive and user friendly Description Deep Learning with Kelp.Net is the ultimate reference for C# .Net developers who are passionate about deep learning. Readers will learn all the skills necessary to develop powerful, scalable and flexible deep learning models from a fluid and easy to use API. Upon completing the book the reader will have all the tools necessary to add powerful deep learning capabilities to their new or existing applications. What will you learn In-depth knowledge of Kelp.Net How to develop deep learning models C# deep learning programming Open-Computing Language (OpenCL) Loading and saving deep learning models How to develop and use activation functions How to test deep learning models Who this book is for This book targets C# .Net developers who are passionate about deep learning yet want to do so from an easy and intuitive API. Table of contents 1. Introduction 2. ML/DL Terms and Concepts 3. Deep Instrumentation 4. Kelp.Net Reference 5. Loading and Saving Models 6. Model Testing and Training 7. Sample Deep Learning Tests 8. Creating Your Own Deep Learning Tests 9. Appendix A: Evaluation Metrics 10. Appendix B: OpenCL About the author Matt R. Cole is a seasoned developer and published author with over 30 years' experience in Microsoft Windows, C, C++, C# and .Net. Matt is the owner of Evolved AI Solutions, a premier provider of advanced Machine Learning/Bio-AI technologies. Matt developed the first enterprise grade MicroService framework written completely in C# and .Net, which is used in production by a major hedge fund in NYC. Matt also developed the first Bio Artificial Intelligence framework which completely integrates mirror and canonical neurons. He continues to push the limits of Machine Learning, Biological Artificial Intelligence, Deep Learning and MicroServices. In his spare time Matt loves to continue his education and contribute to open source efforts such as Kelp.Net. His Website: [www.evolvedaisolutions.com](http://www.evolvedaisolutions.com) His LinkedIn Profile: <https://www.linkedin.com/in/evolvedai/> His Blog:

<https://evolvedaisolutions.com/blog.html>

*Hands-On Machine Learning with ML.NET* Springer

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 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.