Computer Vision And Machine Learning With Rgb D Sensors Advances In Computer Vision And Pattern Recognition

Imagine a world where machines can see and understand the world the way humans do. Rapid progress in artificial intelligence has led to smartphones that recognize faces, cars that detect pedestrians, and algorithms that suggest diagnoses from clinical images, among many other applications. The success of computer vision is founded on a deep understanding of the neural circuits in the brain responsible for visual processing. This book introduces the neuroscientific study of neuronal computations in visual cortex alongside of the psychological understanding of visual cognition and the burgeoning field of biologically-inspired artificial intelligence. Topics include the neurophysiological investigation of visual cortex, visual illusions, visual disorders, deep convolutional neural networks, machine learning, and generative adversarial networks among others. It is an ideal resource for students and researchers looking to build bridges across different approaches to studying and developing visual systems.

The goal of this book is to address the use of several important machine learning techniques into computer vision applications. An innovative combination of computer vision and machine learning techniques has the promise of advancing the field of computer vision, which contributes to better understanding of complex real-world applications. The effective usage of machine learning technology in real-world computer vision problems requires understanding the domain of application, abstraction of a learning problem from a given computer vision task, and the selection of appropriate representations for the learnable (input) and learned (internal) entities of the system. In this book, we address all these important aspects from a new perspective: that the key element in the current computer revolution is the use of machine learning to capture the variations in visual appearance, rather than having the designer of the model accomplish this. As a bonus, models learned from large datasets are likely to be more robust and more realistic than the brittle all-design models.

This book compiles leading research on the development of explainable and interpretable machine learning methods in the context of computer vision and machine learning. Research progress in computer vision and pattern recognition has led to a variety of modeling techniques with almost human-like performance. Although these models have obtained astounding results, they are limited in their explainability and interpretability: what is the rationale behind the decision made? what in the model structure explains its functioning? Hence, while good performance is a critical required characteristic for learning machines, explainability and interpretability capabilities are needed to take learning machines to the next step to include them in decision support systems involving human supervision. This book, written by leading international researchers, addresses key topics of explainability and interpretability, including the following: · Evaluation and Generalization in Interpretable Machine Learning · Explanation Methods in Deep Learning · Learning Functional Causal Models with Generative Neural Networks · Learning Interpreatable Rules for Multi-Label Classification · Structuring Neural Networks for More Explainable Predictions · Generating Post Hoc Rationales of Deep Visual Classification Decisions · Ensembling Visual Explanations · Explainable Deep Driving by Visualizing Causal Attention · Interdisciplinary Perspective on Algorithmic Job Candidate Search · Multimodal Personality Trait Analysis for Explainable Modeling of Job Interview Decisions · Inherent Explainability Pattern Theory-based Video Event Interpretations

Organizations spend huge resources in developing software that can perform the way a human does. Image classification, object detection and tracking, pose estimation, facial recognition, and sentiment estimation all play a major role in solving computer vision problems. This book will bring into focus these and other deep learning architectures and techniques to help you create solutions using Keras and the TensorFlow library. You'll also review mutliple neural network architectures, including LeNet, AlexNet, VGG, Inception, R-CNN, Fast R-CNN, Faster R-CNN, Mask R-CNN, YOLO, and SqueezeNet and see how they work alongside Python code via best practices, tips, tricks, shortcuts, and pitfalls. All code snippets will be broken down and discussed thoroughly so you can implement the same principles in your respective environments. Computer Vision Using Deep Learning offers a comprehensive yet succinct guide that stitches DL and CV together to automate operations, reduce human intervention, increase capability, and cut the costs. What You'll Learn Examine deep learning code and concepts to apply guiding principals to your own projects Classify and evaluate various architectures to better understand your options in various use cases Go behind the scenes of basic deep learning functions to find out how they work Who This Book Is For Professional practitioners working in the fields of software engineering and data science. A working knowledge of Python is strongly recommended. Students and innovators working on advanced degrees in areas related to computer vision and Deep Learning.

If you want a basic understanding of computer vision’s underlying theory and algorithms, this hands-on introduction is the ideal place to start. You’ll learn techniques for object recognition, 3D reconstruction, stereo imaging, augmented reality, and other computer vision applications as you follow clear examples written in Python. Programming Computer Vision with Python explains computer vision in broad terms that won’t bog you down in theory. You get complete code samples with explanations on how to reproduce and build upon each example, along with exercises to help you apply what you’ve learned. This book is ideal for students, researchers, and enthusiasts with basic programming and standard mathematical skills. Learn techniques used in robot navigation, medical image analysis, and other computer vision applications Work with image mappings and transforms, such as texture warping and panorama creation Compute 3D reconstructions from several images of the same scene Organize images based on similarity or content, using clustering methods Build efficient image retrieval techniques to search for images based on visual content Use algorithms to classify image content and recognize objects Access the popular OpenCV library through a Python interface

Using the Pi Camera and a Raspberry Pi board, expand and replicate interesting machine learning (ML) experiments. This book provides a solid overview of ML and a myriad of underlying topics to further explore. Non-technical discussions temper complex technical explanations to make the hottest and most complex topic in the hobbyist world of computing understandable and approachable. Machine learning, also commonly referred to as deep learning (DL), is currently being integrated into a multitude of commercial products as well as widely being used in industrial, medical, and military applications. It is hard to find any modern human activity, which has not been "touched" by artificial intelligence (AI) applications. Building on the concepts first presented in Beginning Artificial Intelligence with the Raspberry Pi, you’ll go beyond simply understanding the concepts of AI into working with real machine learning experiments and applying practical deep learning concepts to experiments with the Pi board and computer vision. What you learn with Machine Learning with the Raspberry Pi can then be moved on to other platforms to go even further in the world of AI and ML to better your hobbyist or commercial projects. What You'll Learn Acquire a working knowledge of current

ML Use the Raspberry Pi to implement ML techniques and algorithms Apply AI and ML tools and techniques to your own work projects and studies Who This Book Is For Engineers and scientists but also experienced makers and hobbyists. Motivated high school students who desire to learn about ML can benefit from this material with determination.

This book will take you through the process of efficiently training deep neural networks in Java for Computer Vision-related tasks. You will build real-world applications ranging from simple Java handwritten digit recognition models to real-time autonomous car driving systems and face recognition models using the popular Java-based libraries.

In this groundbreaking new volume, computer researchers discuss the development of technologies and specific systems that can interpret data with respect to domain knowledge. Although the chapters each illuminate different aspects of image interpretation, all utilize a common approach - one that asserts such interpretation must involve perceptual learning in terms of automated knowledge acquisition and application, as well as feedback and consistency checks between encoding, feature extraction, and the known knowledge structures in a given application domain. The text is profusely illustrated with numerous figures and tables to reinforce the concepts discussed.

How does the computer learn to understand what it sees? Deep Learning for Vision Systems answers that by applying deep learning to computer vision. Using only high school algebra, this book illuminates the concepts behind visual intuition. You'll understand how to use deep learning architectures to build vision system applications for image generation and facial recognition. Summary Computer vision is central to many leading-edge innovations, including self-driving cars, drones, augmented reality, facial recognition, and much, much more. Amazing new computer vision applications are developed every day, thanks to rapid advances in AI and deep learning (DL). Deep Learning for Vision Systems teaches you the concepts and tools for building intelligent, scalable computer vision systems that can identify and react to objects in images, videos, and real life. With author Mohamed Elgendy's expert instruction and illustration of real-world projects, you'll finally grok state-of-the-art deep learning techniques, so you can build, contribute to, and lead in the exciting realm of computer vision! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology How much has computer vision advanced? One ride in a Tesla is the only answer you'll need. Deep learning techniques have led to exciting breakthroughs in facial recognition, interactive simulations, and medical imaging, but nothing beats seeing a car respond to real-world stimuli while speeding down the highway. About the book How does the computer learn to understand what it sees? Deep Learning for Vision Systems answers that by applying deep learning to computer vision. Using only high school algebra, this book illuminates the concepts behind visual intuition. You'll understand how to use deep learning architectures to build vision system applications for image generation and facial recognition. What's inside Image classification and object detection Advanced deep learning architectures Transfer learning and generative adversarial networks DeepDream and neural style transfer Visual embeddings and image search About the reader For intermediate Python programmers. About the author Mohamed Elgendy is the VP of Engineering at Rakuten. A seasoned AI expert, he has previously built and managed AI products at Amazon and Twilio. Table of Contents PART 1 - DEEP LEARNING FOUNDATION 1 Welcome to computer vision 2 Deep learning and neural networks 3 Convolutional neural networks 4 Structuring DL projects and hyperparameter tuning PART 2 - IMAGE CLASSIFICATION AND DETECTION 5 Advanced CNN architectures 6 Transfer learning 7 Object detection with R-CNN, SSD, and YOLO PART 3 - GENERATIVE MODELS AND VISUAL EMBEDDINGS 8 Generative adversarial networks (GANs) 9 DeepDream and neural style transfer 10 Visual embeddings

This book provides a survey of deep learning approaches to domain adaptation in computer vision. It gives the reader an overview of the state-of-the-art research in deep learning based domain adaptation. This book also discusses the various approaches to deep learning based domain adaptation in recent years. It outlines the importance of domain adaptation for the advancement of computer vision, consolidates the research in the area and provides the reader with promising directions for future research in domain adaptation. Divided into four parts, the first part of this book begins with an introduction to domain adaptation, which outlines the problem statement, the role of domain adaptation and the motivation for research in this area. It includes a chapter outlining pre-deep learning era domain adaptation techniques. The second part of this book highlights feature alignment based approaches to domain adaptation. The third part of this book outlines image alignment procedures for domain adaptation. The final section of this book presents novel directions for research in domain adaptation. This book targets researchers working in artificial intelligence, machine learning, deep learning and computer vision. Industry professionals and entrepreneurs seeking to adopt deep learning into their applications will also be interested in this book.

Deep learning algorithms have brought a revolution to the computer vision community by introducing non-traditional and efficient solutions to several image-related problems that had long remained unsolved or partially addressed. This book presents a collection of eleven chapters where each individual chapter explains the deep learning principles of a specific topic, introduces reviews of up-to-date techniques, and presents research findings to the computer vision community. The book covers a broad scope of topics in deep learning concepts and applications such as accelerating the convolutional neural network inference on field-programmable gate arrays, fire detection in surveillance applications, face recognition, action and activity recognition, semantic segmentation for autonomous driving, aerial imagery registration, robot vision, tumor detection, and skin lesion segmentation as well as skin melanoma classification. The content of this book has been organized such that each chapter can be read independently from the others. The book is a valuable companion for researchers, for postgraduate and possibly senior undergraduate students who are taking an advanced course in related topics, and for those who are interested in deep learning with applications in computer vision, image processing, and pattern recognition.

Covariance matrices play important roles in many areas of mathematics, statistics, and machine learning, as well as their applications. In computer vision and image processing, they give rise to a powerful data representation, namely the covariance descriptor, with numerous practical applications. In this book, we begin by presenting an overview of the {\it finite-dimensional covariance matrix} representation approach of images, along with its statistical interpretation. In particular, we discuss the various distances and divergences that arise from the intrinsic geometrical structures of the set of Symmetric Positive Definite (SPD) matrices, namely Riemannian manifold and convex cone structures. Computationally, we focus on kernel methods on covariance matrices, especially using the Log-Euclidean distance. We then show some of the latest developments in the generalization of the finite-dimensional covariance matrix representation to the {\it infinite-dimensional covariance operator} representation via positive definite kernels. We present the generalization of the affine-invariant Riemannian metric and the Log-Hilbert-Schmidt metric, which generalizes the Log Euclidean distance. Computationally, we focus on kernel methods on covariance operators, especially using the Log-Hilbert-Schmidt distance. Specifically, we present a two-layer kernel machine, using the Log-Hilbert-Schmidt distance and its finite-dimensional approximation, which reduces the computational complexity of the exact formulation while largely preserving its capability. Theoretical analysis shows that, mathematically, the approximate Log-Hilbert-Schmidt distance should be preferred over the approximate Log-Hilbert-Schmidt inner product and, computationally, it should be preferred over the approximate affine-invariant Riemannian distance. Numerical experiments on image classification demonstrate significant improvements of the infinite-dimensional formulation over the finite-dimensional counterpart. Given the numerous applications of covariance matrices in many areas of mathematics, statistics, and machine learning, just to name a few, we expect that the infinite-dimensional covariance operator formulation presented here will have many more applications beyond those in computer vision.

Get well versed with state-of-the-art techniques to tailor training processes and boost the performance of computer vision models using

machine learning and deep learning techniques Key Features Develop, train, and use deep learning algorithms for computer vision tasks using TensorFlow 2.x Discover practical recipes to overcome various challenges faced while building computer vision models Enable machines to gain a human level understanding to recognize and analyze digital images and videos Book Description Computer vision is a scientific field that enables machines to identify and process digital images and videos. This book focuses on independent recipes to help you perform various computer vision tasks using TensorFlow. The book begins by taking you through the basics of deep learning for computer vision, along with covering TensorFlow 2.x's key features, such as the Keras and tf.data.Dataset APIs. You'll then learn about the ins and outs of common computer vision tasks, such as image classification, transfer learning, image enhancing and styling, and object detection. The book also covers autoencoders in domains such as inverse image search indexes and image denoising, while offering insights into various architectures used in the recipes, such as convolutional neural networks (CNNs), region-based CNNs (R-CNNs), VGGNet, and You Only Look Once (YOLO). Moving on, you'll discover tips and tricks to solve any problems faced while building various computer vision applications. Finally, you'll delve into more advanced topics such as Generative Adversarial Networks (GANs), video processing, and AutoML, concluding with a section focused on techniques to help you boost the performance of your networks. By the end of this TensorFlow book, you'll be able to confidently tackle a wide range of computer vision problems using TensorFlow 2.x. What you will learn Understand how to detect objects using state-of-the-art models such as YOLOv3 Use AutoML to predict gender and age from images Segment images using different approaches such as FCNs and generative models Learn how to improve your network's performance using rank-N accuracy, label smoothing, and test time augmentation Enable machines to recognize people's emotions in videos and real-time streams Access and reuse advanced TensorFlow Hub models to perform image classification and object detection Generate captions for images using CNNs and RNNs Who this book is for This book is for computer vision developers and engineers, as well as deep learning practitioners looking for go-to solutions to various problems that commonly arise in computer vision. You will discover how to employ modern machine learning (ML) techniques and deep learning architectures to perform a plethora of computer vision tasks. Basic knowledge of Python programming and computer vision is required.

By using machine learning models to extract information from images, organizations today are making breakthroughs in healthcare, manufacturing, retail, and other industries. This practical book shows ML engineers and data scientists how to solve a variety of image problems including classification, object detection, autoencoders, image generation, counting, and captioning with proven ML techniques. Google engineers Valliappa Lakshmanan, Martin Garner, and Ryan Gillard show you how to develop accurate and explainable computer vision ML models and put them into large-scale production using robust ML architecture in a flexible and maintainable way. You'll learn how to design, train, evaluate, and predict with models written in TensorFlow/Keras. This book also covers best practices to improve the operationalization of the models using end-to-end ML pipelines. You'll learn how to: Design ML architecture for computer vision tasks Select a model (such as ResNet, SqueezeNet, or EfficientNet) appropriate to your task Create an end-to-end ML pipeline to train, evaluate, deploy, and explain your model Preprocess images for data augmentation and to support learnability Incorporate explainability and responsible AI best practices Deploy image models as web services or on edge devices Monitor and manage ML models

This practical book shows you how to employ machine learning models to extract information from images. ML engineers and data scientists will learn how to solve a variety of image problems including classification, object detection, autoencoders, image generation, counting, and captioning with proven ML techniques. This book provides a great introduction to end-to-end deep learning: dataset creation, data preprocessing, model design, model training, evaluation, deployment, and interpretability. Google engineers Valliappa Lakshmanan, Martin Görner, and Ryan Gillard show you how to develop accurate and explainable computer vision ML models and put them into large-scale production using robust ML architecture in a flexible and maintainable way. You'll learn how to design, train, evaluate, and predict with models written in TensorFlow or Keras. You'll learn how to: Design ML architecture for computer vision tasks Select a model (such as ResNet, SqueezeNet, or EfficientNet) appropriate to your task Create an end-to-end ML pipeline to train, evaluate, deploy, and explain your model Preprocess images for data augmentation and to support learnability Incorporate explainability and responsible AI best practices Deploy image models as web services or on edge devices Monitor and manage ML models

This book discusses computer vision, a noncontact as well as a nondestructive technique involving the development of theoretical and algorithmic tools for automatic visual understanding and recognition which finds huge applications in agricultural productions. It also entails how rendering of machine learning techniques to computer vision algorithms is boosting this sector with better productivity by developing more precise systems. Computer vision and machine learning (CV-ML) helps in plant disease assessment along with crop condition monitoring to control the degradation of yield, quality, and severe financial loss for farmers. Significant scientific and technological advances have been made in defect assessment, quality grading, disease recognition, pests, insects, fruits, and vegetable types recognition and evaluation of a wide range of agricultural plants, crops, leaves, and fruits. The book discusses intelligent robots developed with the touch of CV-ML which can help farmers to perform various tasks like planting, weeding, harvesting, plant health monitoring, and so on. The topics covered in the book include plant, leaf, and fruit disease detection, crop health monitoring, applications of robots in agriculture, precision farming, assessment of product quality and defects, pest, insect, fruits, and vegetable types recognition.

Computer Vision: Principles, Algorithms, Applications, Learning (previously entitled Computer and Machine Vision) clearly and systematically presents the basic methodology of computer vision, covering the essential elements of the theory while emphasizing algorithmic and practical design constraints. This fully revised fifth edition has brought in more of the concepts and applications of computer vision, making it a very comprehensive and up-to-date text suitable for undergraduate and graduate students, researchers and R&D engineers working in this vibrant subject. See an interview with the author explaining his approach to teaching and learning computer vision - http://scitechconnect.elsevier.com/computer-vision/ Three new chapters on Machine Learning emphasise the way the subject has been developing; Two chapters cover Basic Classification Concepts and Probabilistic Models; and the The third covers the principles of Deep Learning Networks and shows their impact on computer vision, reflected in a new chapter Face Detection and Recognition. A new chapter on Object Segmentation and Shape Models reflects the methodology of machine learning and gives practical demonstrations of its application. In-depth discussions have been included on geometric transformations, the EM algorithm, boosting, semantic segmentation, face frontalisation, RNNs and other key topics. Examples and applications—including the location of biscuits, foreign bodies, faces, eyes, road lanes, surveillance, vehicles and pedestrians—give the ‘ins and outs’ of developing real-world vision systems, showing the realities of practical implementation. Necessary mathematics and essential theory are made approachable by careful explanations and well-illustrated examples. The ‘recent developments’ sections included in each chapter aim to bring students and practitioners up to date with this fast-moving subject. Tailored programming examples—code, methods, illustrations, tasks, hints and solutions (mainly involving MATLAB and C++)

Apply neural network architectures to build state-of-the-art computer vision applications using the Python programming language Key Features Gain a fundamental understanding of advanced computer vision and neural network models in use today Cover tasks such as low-level vision, image classification, and object detection Develop deep learning models on cloud platforms and optimize them using TensorFlow Lite and the OpenVINO toolkit Book Description Computer vision allows machines to gain human-level understanding to visualize, process, and analyze images and videos. This book focuses on using TensorFlow to help you learn advanced computer vision tasks such as image acquisition, processing, and analysis. You'll start with the key principles of computer vision and deep learning to build a solid foundation, before covering neural network architectures and understanding how they work rather than using them as a black box. Next, you'll explore

architectures such as VGG, ResNet, Inception, R-CNN, SSD, YOLO, and MobileNet. As you advance, you'll learn to use visual search methods using transfer learning. You'll also cover advanced computer vision concepts such as semantic segmentation, image inpainting with GAN's, object tracking, video segmentation, and action recognition. Later, the book focuses on how machine learning and deep learning concepts can be used to perform tasks such as edge detection and face recognition. You'll then discover how to develop powerful neural network models on your PC and on various cloud platforms. Finally, you'll learn to perform model optimization methods to deploy models on edge devices for real-time inference. By the end of this book, you'll have a solid understanding of computer vision and be able to confidently develop models to automate tasks. What you will learn Explore methods of feature extraction and image retrieval and visualize different layers of the neural network model Use TensorFlow for various visual search methods for real-world scenarios Build neural networks or adjust parameters to optimize the performance of models Understand TensorFlow DeepLab to perform semantic segmentation on images and DCGAN for image inpainting Evaluate your model and optimize and integrate it into your application to operate at scale Get up to speed with techniques for performing manual and automated image annotation Who this book is for This book is for computer vision professionals, image processing professionals, machine learning engineers and AI developers who have some knowledge of machine learning and deep learning and want to build expert-level computer vision applications. In addition to familiarity with TensorFlow, Python knowledge will be required to get started with this book.

Numerical Algorithms: Methods for Computer Vision, Machine Learning, and Graphics presents a new approach to numerical analysis for modern computer scientists. Using examples from a broad base of computational tasks, including data processing, computational photography, and animation, the textbook introduces numerical modeling and algorithmic desig

Machine learning allows for non-conventional and productive answers for issues within various fields, including problems related to visually perceptive computers. Applying these strategies and algorithms to the area of computer vision allows for higher achievement in tasks such as spatial recognition, big data collection, and image processing. There is a need for research that seeks to understand the development and efficiency of current methods that enable machines to see. Challenges and Applications for Implementing Machine Learning in Computer Vision is a collection of innovative research that combines theory and practice on adopting the latest deep learning advancements for machines capable of visual processing. Highlighting a wide range of topics such as video segmentation, object recognition, and 3D modelling, this publication is ideally designed for computer scientists, medical professionals, computer engineers, information technology practitioners, industry experts, scholars, researchers, and students seeking current research on the utilization of evolving computer vision techniques.

Learn how to model and train advanced neural networks to implement a variety of Computer Vision tasks Key Features Train different kinds of deep learning model from scratch to solve specific problems in Computer Vision Combine the power of Python, Keras, and TensorFlow to build deep learning models for object detection, image classification, similarity learning, image captioning, and more Includes tips on optimizing and improving the performance of your models under various constraints Book Description Deep learning has shown its power in several application areas of Artificial Intelligence, especially in Computer Vision. Computer Vision is the science of understanding and manipulating images, and finds enormous applications in the areas of robotics, automation, and so on. This book will also show you, with practical examples, how to develop Computer Vision applications by leveraging the power of deep learning. In this book, you will learn different techniques related to object classification, object detection, image segmentation, captioning, image generation, face analysis, and more. You will also explore their applications using popular Python libraries such as TensorFlow and Keras. This book will help you master state-of-the-art, deep learning algorithms and their implementation. What you will learn Set up an environment for deep learning with Python, TensorFlow, and Keras Define and train a model for image and video classification Use features from a pre-trained Convolutional Neural Network model for image retrieval Understand and implement object detection using the real-world Pedestrian Detection scenario Learn about various problems in image captioning and how to overcome them by training images and text together Implement similarity matching and train a model for face recognition Understand the concept of generative models and use them for image generation Deploy your deep learning models and optimize them for high performance Who this book is for This book is targeted at data scientists and Computer Vision practitioners who wish to apply the concepts of Deep Learning to overcome any problem related to Computer Vision. A basic knowledge of programming in Python—and some understanding of machine learning concepts—is required to get the best out of this book.

Computer vision is the science and technology of making machines that see. It is concerned with the theory, design and implementation of algorithms that can automatically process visual data to recognize objects, track and recover their shape and spatial layout. The International Computer Vision Summer School - ICVSS was established in 2007 to provide both an objective and clear overview and an in-depth analysis of the state-of-the-art research in Computer Vision. The courses are delivered by world renowned experts in the field, from both academia and industry, and cover both theoretical and practical aspects of real Computer Vision problems. The school is organized every year by University of Cambridge (Computer Vision and Robotics Group) and University of Catania (Image Processing Lab). Different topics are covered each year. A summary of the past Computer Vision Summer Schools can be found at: http://www.dmi.unict.it/icvss This edited volume contains a selection of articles covering some of the talks and tutorials held during the last editions of the school. The chapters provide an in-depth overview of challenging areas with key references to the existing literature.

Build practical applications of computer vision using the OpenCV library with Python. This book discusses different facets of computer vision such as image and object detection, tracking and motion analysis and their applications with examples. The author starts with an introduction to computer vision followed by setting up OpenCV from scratch using Python. The next section discusses specialized image processing and segmentation and how images are stored and processed by a computer. This involves pattern recognition and image tagging using the OpenCV library. Next, you'll work with object detection, video storage and interpretation, and human detection using OpenCV. Tracking and motion is also discussed in detail. The book also discusses creating complex deep learning models with CNN and RNN. The author finally concludes with recent applications and trends in computer vision. After reading this book, you will be able to understand and implement computer vision and its applications with OpenCV using Python. You will also be able to create deep learning models with CNN and RNN and understand how these cutting-edge deep learning architectures work. What You Will Learn Understand what computer vision is, and its overall application in intelligent automation systems Discover the deep learning techniques required to build computer vision applications Build complex computer vision applications using the latest techniques in OpenCV, Python, and NumPy Create practical applications and implementations such as face detection and recognition, handwriting recognition, object detection, and tracking and motion analysis Who This Book Is ForThose who have a basic understanding of machine learning and Python and are looking to learn computer vision and its applications.

Computer vision is achieving a new frontier of capabilities in fields like health, automobile or robotics. This book explores TensorFlow 2, Google's open-source AI framework, and teaches how to leverage deep neural networks for visual tasks. It will help you acquire the insight and skills to be a part of the exciting advances in computer vision.

Deploy deep learning applications into production across multiple platforms. You will work on computer vision applications that use the convolutional neural network (CNN) deep learning model and Python. This book starts by explaining the traditional machine-learning pipeline, where you will analyze an image dataset. Along the way you will cover artificial neural networks (ANNs), building one from scratch in Python, before optimizing it using genetic algorithms.

For automating the process, the book highlights the limitations of traditional hand-crafted features for computer vision and why the CNN deep-learning model is the state-of-art solution. CNNs are discussed from scratch to demonstrate how they are different and more efficient than the fully connected ANN (FCNN). You will implement a CNN in Python to give you a full understanding of the model. After consolidating the basics, you will use TensorFlow to build a practical image-recognition model that you will deploy to a web server using Flask, making it accessible over the Internet. Using Kivy and NumPy, you will create cross-platform data science applications with low overheads. This book will help you apply deep learning and computer vision concepts from scratch, step-by-step from conception to production. What You Will Learn Understand how ANNs and CNNs work Create computer vision applications and CNNs from scratch using Python Follow a deep learning project from conception to production using TensorFlow Use NumPy with Kivy to build cross-platform data science applications Who This Book Is ForData scientists, machine learning and deep learning engineers, software developers.

Starting from the basics of neural networks, this book covers over 50 applications of computer vision and helps you to gain a solid understanding of the theory of various architectures before implementing them. Each use case is accompanied by a notebook in GitHub with ready-to-execute code and self-assessment questions.

Updated for OpenCV 4 and Python 3, this book covers the latest on depth cameras, 3D tracking, augmented reality, and deep neural networks, helping you solve real-world computer vision problems with practical code Key Features Build powerful computer vision applications in concise code with OpenCV 4 and Python 3 Learn the fundamental concepts of image processing, object classification, and 2D and 3D tracking Train, use, and understand machine learning models such as Support Vector Machines (SVMs) and neural networks Book Description Computer vision is a rapidly evolving science, encompassing diverse applications and techniques. This book will not only help those who are getting started with computer vision but also experts in the domain. You'll be able to put theory into practice by building apps with OpenCV 4 and Python 3. You'll start by understanding OpenCV 4 and how to set it up with Python 3 on various platforms. Next, you'll learn how to perform basic operations such as reading, writing, manipulating, and displaying still images, videos, and camera feeds. From taking you through image processing, video analysis, and depth estimation and segmentation, to helping you gain practice by building a GUI app, this book ensures you'll have opportunities for hands-on activities. Next, you'll tackle two popular challenges: face detection and face recognition. You'll also learn about object classification and machine learning concepts, which will enable you to create and use object detectors and classifiers, and even track objects in movies or video camera feed. Later, you'll develop your skills in 3D tracking and augmented reality. Finally, you'll cover ANNs and DNNs, learning how to develop apps for recognizing handwritten digits and classifying a person's gender and age. By the end of this book, you'll have the skills you need to execute real-world computer vision projects. What you will learn Install and familiarize yourself with OpenCV 4's Python 3 bindings Understand image processing and video analysis basics Use a depth camera to distinguish foreground and background regions Detect and identify objects, and track their motion in videos Train and use your own models to match images and classify objects Detect and recognize faces, and classify their gender and age Build an augmented reality application to track an image in 3D Work with machine learning models, including SVMs, artificial neural networks (ANNs), and deep neural networks (DNNs) Who this book is for If you are interested in learning computer vision, machine learning, and OpenCV in the context of practical real-world applications, then this book is for you. This OpenCV book will also be useful for anyone getting started with computer vision as well as experts who want to stay up-to-date with OpenCV 4 and Python 3. Although no prior knowledge of image processing, computer vision or machine learning is required, familiarity with basic Python programming is a must.

Written by a team of International experts, this edited book covers state-of-the-art research in the fields of computer vision and recognition systems from fundamental concepts to methodologies and technologies and real-world applications. The book will be useful for industry and academic researchers, scientists and engineers.

A modern treatment focusing on learning and inference, with minimal prerequisites, real-world examples and implementable algorithms.

Advanced Methods and Deep Learning in Computer Vision presents advanced computer vision methods, emphasizing machine and deep learning techniques that have emerged during the past 5–10 years. The book provides clear explanations of principles and algorithms supported with applications. Topics covered include machine learning, deep learning networks, generative adversarial networks, deep reinforcement learning, self-supervised learning, extraction of robust features, object detection, semantic segmentation, linguistic descriptions of images, visual search, visual tracking, 3D shape retrieval, image inpainting, novelty and anomaly detection. This book provides easy learning for researchers and practitioners of advanced computer vision methods, but it is also suitable as a textbook for a second course on computer vision and deep learning for advanced undergraduates and graduate students. Provides an important reference on deep learning and advanced computer methods that was created by leaders in the field Illustrates principles with modern, real-world applications Suitable for self-learning or as a text for graduate courses

Computer Vision: Algorithms and Applications explores the variety of techniques commonly used to analyze and interpret images. It also describes challenging real-world applications where vision is being successfully used, both for specialized applications such as medical imaging, and for fun, consumer-level tasks such as image editing and stitching, which students can apply to their own personal photos and videos. More than just a source of "recipes," this exceptionally authoritative and comprehensive textbook/reference also takes a scientific approach to basic vision problems, formulating physical models of the imaging process before inverting them to produce descriptions of a scene. These problems are also analyzed using statistical models and solved using rigorous engineering techniques. Topics and features: structured to support active curricula and project-oriented courses, with tips in the Introduction for using the book in a variety of

customized courses; presents exercises at the end of each chapter with a heavy emphasis on testing algorithms and containing numerous suggestions for small mid-term projects; provides additional material and more detailed mathematical topics in the Appendices, which cover linear algebra, numerical techniques, and Bayesian estimation theory; suggests additional reading at the end of each chapter, including the latest research in each sub-field, in addition to a full Bibliography at the end of the book; supplies supplementary course material for students at the associated website, http://szeliski.org/Book/. Suitable for an upper-level undergraduate or graduate-level course in computer science or engineering, this textbook focuses on basic techniques that work under real-world conditions and encourages students to push their creative boundaries. Its design and exposition also make it eminently suitable as a unique reference to the fundamental techniques and current research literature in computer vision.

This book presents an interdisciplinary selection of cutting-edge research on RGB-D based computer vision. Features: discusses the calibration of color and depth cameras, the reduction of noise on depth maps and methods for capturing human performance in 3D; reviews a selection of applications which use RGB-D information to reconstruct human figures, evaluate energy consumption and obtain accurate action classification; presents an approach for 3D object retrieval and for the reconstruction of gas flow from multiple Kinect cameras; describes an RGB-D computer vision system designed to assist the visually impaired and another for smart-environment sensing to assist elderly and disabled people; examines the effective features that characterize static hand poses and introduces a unified framework to enforce both temporal and spatial constraints for hand parsing; proposes a new classifier architecture for real-time hand pose recognition and a novel hand segmentation and gesture recognition system.

"This book provides a working guide to the C++ Open Source Computer Vision Library (OpenCV) version 3.x and gives a general background on the field of computer vision sufficient to help readers use OpenCV effectively."--Preface.

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Practical Machine Learning for Computer Vision"O'Reilly Media, Inc."

"Over the last years, kernel methods have established themselves as powerful tools for computer vision researchers as well as for practitioners. In this tutorial, we give an introduction to kernel methods in computer vision from a geometric perspective, introducing not only the ubiquitous support vector machines, but also less known techniques for regression, dimensionality reduction, outlier detection, and clustering. Additionally, we give an outlook on very recent, non-classical techniques for the prediction of structure data, for the estimation of statistical dependency, and for learning the kernel function itself. All methods are illustrated with examples of successful application from the recent computer vision research literature" --Abstract.

With The Computer Vision Workshop, you’ll explore the basic and advanced techniques in video and image processing using OpenCV and Python. It is filled with real-world exercises and activities that will make the learning process easy and enjoyable.

Step-by-step tutorials on deep learning neural networks for computer vision in python with Keras.