
Computer Vision And Machine Learning With Rgb D Sensors

Advances In Computer Vision And Pattern Recognition

Computer Vision

Computer Vision and Machine Learning with RGB-D Sensors

Implement machine learning solutions to overcome various computer vision challenges

Practical Deep Learning for Cloud, Mobile, and Edge

Computer Vision

Methods for Computer Vision, Machine Learning, and Graphics

Deep Learning in Computer Vision

Deep Learning in Computer Vision

Modern Computer Vision with PyTorch

Computer Vision and Recognition Systems Using Machine and Deep Learning Approaches

A Perceptual Organization Approach to Computer Vision and Machine Learning

Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow

Deep Learning for Vision Systems

Mastering Computer Vision with TensorFlow 2.x

With Detailed Examples in Python Using TensorFlow and Kivy

Elements of Deep Learning for Computer Vision

Principles, Algorithms, Applications, Learning

Principles and Applications

Hands-On Computer Vision with TensorFlow 2

Leverage deep learning to create powerful image processing apps with TensorFlow 2.0 and Keras

Expert techniques to train advanced neural networks using TensorFlow and Keras

Build advanced computer vision applications using machine learning and deep learning techniques

Building Computer Vision Projects with OpenCV 4 and C++

Domain Adaptation in Computer Vision with Deep Learning

An Introduction

Computer Vision Using Deep Learning

Machine Learning with the Raspberry Pi
Programming Computer Vision with Python
Practical Machine Learning for Computer Vision
Neural Network Architectures with Python and Keras
Get to grips with tools, techniques, and algorithms for computer vision and machine learning, 3rd Edition
Building Computer Vision Applications Using Artificial Neural Networks
With Step-by-Step Examples in OpenCV and TensorFlow with Python
Improve the Efficiency of Artificial Intelligence
Challenges and Applications for Implementing Machine Learning in Computer Vision
Deep Learning for Computer Vision with SAS
Practical Machine Learning for Computer Vision
Tools and algorithms for analyzing images
Tensor Voting

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

*Downloaded from
archive.imba.com by
guest*

MELINA ALVARADO

Computer Vision "O'Reilly Media, Inc."
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 design

Computer Vision and Machine Learning with RGB-D Sensors Apress 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 multiple 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.

Implement machine learning solutions to overcome various computer vision challenges Manning Publications 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 For Data scientists, machine learning and deep learning engineers, software developers.

Practical Deep Learning for Cloud, Mobile, and Edge CRC Press
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.

Computer Vision IGI Global
Practical Machine Learning for Computer Vision O'Reilly Media, Inc."
Methods for Computer Vision, Machine Learning, and Graphics Packt Publishing Ltd

Apply computer vision and machine learning concepts in developing business

and industrial applications using a practical, step-by-step approach. The book comprises four main sections starting with setting up your programming environment and configuring your computer with all the prerequisites to run the code examples. Section 1 covers the basics of image and video processing with code examples of how to manipulate and extract useful information from the images. You will mainly use OpenCV with Python to work with examples in this section. Section 2 describes machine learning and neural network concepts as applied to computer vision. You will learn different algorithms of the neural network, such as convolutional neural network (CNN), region-based convolutional neural network (R-CNN), and YOLO. In this

section, you will also learn how to train, tune, and manage neural networks for computer vision. Section 3 provides step-by-step examples of developing business and industrial applications, such as facial recognition in video surveillance and surface defect detection in manufacturing. The final section is about training neural networks involving a large number of images on cloud infrastructure, such as Amazon AWS, Google Cloud Platform, and Microsoft Azure. It walks you through the process of training distributed neural networks for computer vision on GPU-based cloud infrastructure. By the time you finish reading *Building Computer Vision Applications Using Artificial Neural Networks* and working through the code examples, you will have developed some

real-world use cases of computer vision with deep learning. What You Will Learn · Employ image processing, manipulation, and feature extraction techniques · Work with various deep learning algorithms for computer vision · Train, manage, and tune hyperparameters of CNNs and object detection models, such as R-CNN, SSD, and YOLO · Build neural network models using Keras and TensorFlow · Discover best practices when implementing computer vision applications in business and industry · Train distributed models on GPU-based cloud infrastructure Who This Book Is For Data scientists, analysts, and machine learning and software engineering professionals with Python programming knowledge. [Deep Learning in Computer Vision](#) Packt

Publishing Ltd

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. *Deep Learning in Computer Vision*
Machine Learning Mastery

Discover deep learning and computer vision with SAS! Deep Learning for Computer Vision with SAS®: An Introduction introduces the pivotal components of deep learning. Readers will gain an in-depth understanding of how to build deep feedforward and convolutional neural networks, as well as variants of denoising autoencoders. Transfer learning is covered to help readers learn about this emerging field. Containing a mix of theory and application, this book will also briefly cover methods for customizing deep learning models to solve novel business problems or answer research questions. SAS programs and data are included to reinforce key concepts and allow readers to follow along with included demonstrations. Readers will learn how

to: Define and understand deep learning Build models using deep learning techniques and SAS Viya Apply models to score (inference) new data Modify data for better analysis results Search the hyperparameter space of a deep learning model Leverage transfer learning using supervised and unsupervised methods

Modern Computer Vision with PyTorch

Packt Publishing Ltd

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.

Computer Vision and Recognition Systems Using Machine and Deep Learning Approaches Apress

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.

A Perceptual Organization Approach to Computer Vision and Machine Learning

Packt Publishing Ltd

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.

Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow Cambridge University Press
A modern treatment focusing on learning and inference, with minimal prerequisites, real-world examples and implementable algorithms.

Deep Learning for Vision Systems
Springer Nature

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.

Mastering Computer Vision with

TensorFlow 2.x Springer Nature

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

With Detailed Examples in Python Using TensorFlow and Kivy BPB Publications

Conceptualizing deep learning in computer vision applications using PyTorch and Python libraries. **KEY FEATURES** ● Covers a variety of computer vision projects, including face recognition and object recognition such as Yolo, Faster R-CNN. ● Includes graphical representations and illustrations of neural networks and teaches how to program them. ● Includes deep learning techniques and architectures introduced by Microsoft, Google, and the University of Oxford.

DESCRIPTION Elements of Deep Learning for Computer Vision gives a thorough understanding of deep learning and provides highly accurate computer vision solutions while using libraries like PyTorch. This book introduces you to Deep Learning and explains all the concepts required to understand the basic working, development, and tuning of a neural network using Pytorch. The book then addresses the field of computer vision using two libraries, including the Python wrapper/version of OpenCV and PIL. After establishing and understanding both the primary concepts, the book addresses them together by explaining Convolutional Neural Networks(CNNs). CNNs are further elaborated using top industry standards and research to explain how

they provide complicated Object Detection in images and videos, while also explaining their evaluation. Towards the end, the book explains how to develop a fully functional object detection model, including its deployment over APIs. By the end of this book, you are well-equipped with the role of deep learning in the field of computer vision along with a guided process to design deep learning solutions. WHAT YOU WILL LEARN ● Get to know the mechanism of deep learning and how neural networks operate. ● Learn to develop a highly accurate neural network model. ● Access to rich Python libraries to address computer vision challenges. ● Build deep learning models using PyTorch and learn how to deploy using the API. ● Learn to develop

Object Detection and Face Recognition models along with their deployment. WHO THIS BOOK IS FOR This book is for the readers who aspire to gain a strong fundamental understanding of how to infuse deep learning into computer vision and image processing applications. Readers are expected to have intermediate Python skills. No previous knowledge of PyTorch and Computer Vision is required. TABLE OF CONTENTS 1. An Introduction to Deep Learning 2. Supervised Learning 3. Gradient Descent 4. OpenCV with Python 5. Python Imaging Library and Pillow 6. Introduction to Convolutional Neural Networks 7. GoogLeNet, VGGNet, and ResNet 8. Understanding Object Detection 9. Popular Algorithms for Object Detection 10. Faster RCNN with

PyTorch and YoloV4 with Darknet 11.
Comparing Algorithms and API
Deployment with Flask 12. Applications
in Real World
Elements of Deep Learning for Computer
Vision Academic Press
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++)
Principles, Algorithms, Applications,
Learning CRC Press

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.

Principles and Applications "O'Reilly Media, Inc."

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.

Hands-On Computer Vision with TensorFlow 2 Practical Machine Learning for Computer Vision

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
[Leverage deep learning to create powerful image processing apps with](#)

[TensorFlow 2.0 and Keras](#) "O'Reilly Media, Inc."

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.

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

- Sedition Act Of 1918 Definition Us History : [click here](#)