

Concise Computer Vision An Introduction Into Theory And Algorithms Undergraduate Topics In Computer Science

This easy-to-follow textbook presents an engaging introduction to the fascinating world of medical image analysis. Avoiding an overly mathematical treatment, the text focuses on intuitive explanations, illustrating the key algorithms and concepts in a way which will make sense to students from a broad range of different backgrounds. Topics and features: explains what light is, and how it can be captured by a camera and converted into an image, as well as how images can be compressed and stored; describes basic image manipulation methods for understanding and improving image quality, and a useful segmentation algorithm; reviews the basic image processing methods for segmenting or enhancing certain features in an image, with a focus on morphology methods for binary images; examines how to detect, describe, and recognize objects in an image, and how the nature of color can be used for segmenting objects; introduces a statistical method to determine what class of object the pixels in an image represent; describes how to change the geometry within an image, how to align two images so that they are as similar as possible, and how to detect lines and paths in images; provides further exercises and other supplementary material at an associated website. This concise and accessible textbook will be invaluable to undergraduate students of computer science, engineering, medicine, and any multi-disciplinary courses that combine topics on health with data science. Medical practitioners working with medical imaging devices will also appreciate this easy-to-understand explanation of the technology. Based on the highly successful 3-volume reference Handbook of Computer Vision and Applications, this concise edition covers in a single volume the entire spectrum of computer vision ranging from the imaging process to high-end algorithms and applications. This book consists of three parts, including an application gallery. Bridges the gap between theory and practical applications Covers modern concepts in computer vision as well as modern developments in imaging sensor technology Presents a unique interdisciplinary approach covering different areas of modern science

This textbook presents the fundamental concepts and methods for understanding and working with images and video in a unique, easy-to-read style which ensures the material is accessible to a wide audience. Exploring more than just the basics of image processing, the text provides a specific focus on the practical design and implementation of real systems for processing video data. Features: includes more than 100 exercises, as well as C-code snippets of the key algorithms; covers topics on image acquisition, color images, point processing, neighborhood processing, morphology, BLOB analysis, segmentation in video, tracking, geometric transformation, and visual effects; requires only a minimal understanding of mathematics; presents two chapters dedicated to applications; provides a guide to defining suitable values for parameters in video and image processing systems, and to conversion between the RGB color representation and the HIS, HSV and YUV/YCbCr color representations.

This book is suitable for use in a university-level first course in computing (CS1), as well as the increasingly popular course known as CS0. It is difficult for many students to master basic concepts in computer science and programming. A large portion of the confusion can be blamed on the complexity of the tools and materials that are traditionally used to teach CS1 and CS2. This textbook was written with a single overarching goal: to present the core concepts of computer science as simply as possible without being simplistic.

This book addresses one of the most important unsolved problems in artificial intelligence: the task of learning, in an unsupervised manner, from massive quantities of spatiotemporal visual data that are available at low cost. The book covers important scientific discoveries and findings, with a focus on the latest advances in the field. Presenting a coherent structure, the book logically connects novel mathematical formulations and efficient computational solutions for a range of unsupervised learning tasks, including visual feature matching, learning and classification, object discovery, and semantic segmentation in video. The final part of the book proposes a general strategy for visual learning over several generations of student-teacher neural networks, along with a unique view on the future of unsupervised learning in real-world contexts. Offering a fresh approach to this difficult problem, several efficient, state-of-the-art unsupervised learning algorithms are reviewed in detail, complete with an analysis of their performance on various tasks, datasets, and experimental setups. By highlighting the interconnections between these methods, many seemingly diverse problems are elegantly brought together in a unified way. Serving as an invaluable guide to the computational tools and algorithms required to tackle the exciting challenges in the field, this book is a must-read for graduate students seeking a greater understanding of unsupervised learning, as well as researchers in computer vision, machine learning, robotics, and related disciplines.

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Computer vision encompasses the construction of integrated vision systems and the application of vision to problems of real-world importance. The process of creating 3D models is still rather difficult, requiring mechanical measurement of the camera positions or manual alignment of partial 3D views of a scene. However using algorithms, it is possible to take a collection of stereo-pair images of a scene and then automatically produce a photo-realistic, geometrically accurate digital 3D model. This book provides a comprehensive introduction to the methods, theories and algorithms of 3D computer vision. Almost every theoretical issue is underpinned with practical implementation or a working algorithm using pseudo-code and complete code written in C++ and MatLab®. There is the additional clarification of an accompanying website with downloadable software, case studies and exercises. Organised in three parts, Cyganek and Siebert give a brief history of vision research, and subsequently: present basic low-level image processing operations for image matching, including a separate chapter on image matching algorithms; explain scale-space vision, as well as space reconstruction and multiview integration; demonstrate a variety of practical applications for 3D surface imaging and analysis; provide concise appendices on topics such as the basics of projective geometry and tensor calculus for image processing, distortion and noise in images plus image warping procedures. An Introduction to 3D Computer Vision Algorithms and Techniques is a valuable reference for practitioners and programmers working in 3D computer vision, image processing and analysis as well as computer visualisation. It would also be of interest to advanced students and researchers in the fields of engineering, computer science, clinical photography, robotics, graphics and mathematics.

The fields of computer vision and image processing are constantly evolving as new research and applications in these areas emerge. Staying abreast of the most up-to-date developments in this field is necessary in order to promote further research and apply these developments in real-world settings. Computer Vision: Concepts, Methodologies, Tools, and Applications is an innovative reference source for the latest academic material on development of computers for gaining understanding about videos and digital images. Highlighting a range of topics, such as computational models, machine learning, and image processing, this multi-volume book is ideally designed for academicians, technology professionals, students, and researchers interested in uncovering the latest innovations in the field.

Image recognition has become an increasingly dynamic field with new and emerging civil and military applications in security, exploration, and robotics. Written by experts in fractal-based image and video compression, *A Concise Introduction to Image Processing using C++* strengthens your knowledge of fundamental principles in image acquisition, con
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.

Using a progressive intuitive/mathematical approach, this introduction to computer vision provides necessary theory and examples for practitioners who work in fields where significant information must be extracted automatically from images-- including those interested in multimedia, art and design, geographic information systems, and image databases, in addition to the traditional areas of automation, image science, medical imaging, remote sensing and computer cartography. The book provides a basic set of fundamental concepts, (representations of image information, extraction of 3D scene information from 2D images, etc.) algorithms for analyzing images, and discusses some of the exciting evolving application areas of computer vision. The approach is language and software independent, and includes two significant commercial case studies. Imaging and Image Representation. Binary Image Analysis. Pattern Recognition Concepts. Filtering and Enhancing Images. Color and Shading. Texture. Content-Based Image Retrieval. Motion from 2D Image Sequences. Image Segmentation. Matching in 2D. Perceiving 3D from 2D Images. 3D Sensing and Object Pose Computation. 3D Models and Matching. Virtual Reality. Case Studies. For practitioners in any field where information must be extracted automatically from images.

"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.

[FIRST EDITION] This accessible textbook presents an introduction to computer vision algorithms for industrially-relevant applications of X-ray testing. Features: introduces the mathematical background for monocular and multiple view geometry; describes the main techniques for image processing used in X-ray testing; presents a range of different representations for X-ray images, explaining how these enable new features to be extracted from the original image; examines a range of known X-ray image classifiers and classification strategies; discusses some basic concepts for the simulation of X-ray images and presents simple geometric and imaging models that can be used in the simulation; reviews a variety of applications for X-ray testing, from industrial inspection and baggage screening to the quality control of natural products; provides supporting material at an associated website, including a database of X-ray images and a Matlab toolbox for use with the book's many examples.

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++)
Appropriate for upper-division undergraduate and graduate level courses in computer vision found in departments of computer science, computer engineering and electrical engineering, this book offers a treatment of modern computer vision methods.

Readers discover a contemporary treatment of image processing that balances a broad coverage of major subject areas with in-depth examination of the most foundational topics. *IMAGE PROCESSING AND ANALYSIS* offers an accessible presentation that provides higher-level discussions to challenge the most advanced readers. The book effectively balances key topics from the field of image processing in a format that gradually progresses from easy to more challenging material, while consistently reinforcing a fundamental understanding of the core concepts. The book's hands-on learning approach and full-color presentation allows readers to begin working with images immediately. The book encourages programming as it incorporates algorithmic details and hints, using detailed pseudocode to facilitate an understanding of algorithms and aid in implementation. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

This gentle introduction to High Performance Computing (HPC) for Data Science using the Message Passing Interface (MPI) standard has been designed as a first course for undergraduates on parallel programming on distributed memory models, and requires only basic programming notions. Divided into two parts the first part covers high performance computing using C++ with the Message Passing Interface (MPI) standard followed by a second part providing high-performance data analytics on computer clusters. In the first part, the fundamental notions of blocking versus non-blocking point-to-point communications, global communications (like broadcast or scatter) and collaborative computations (reduce), with Amdalh and Gustafson speed-up laws are described before addressing parallel sorting and parallel linear algebra on computer clusters. The common ring, torus and hypercube topologies of clusters are then explained and global communication procedures on these topologies are studied. This first part closes with the MapReduce (MR) model of computation well-suited to processing big data using the MPI framework. In the second part, the book focuses on high-performance data analytics. Flat and hierarchical clustering algorithms are introduced for data exploration along with how to program these algorithms on computer clusters, followed by machine learning classification, and an introduction to graph analytics. This part closes with a concise introduction to data core-sets that let big data problems be amenable to tiny data problems. Exercises are included at the end of each chapter in order for students to practice the concepts learned, and a final section contains an overall exam which allows them to evaluate how well they have assimilated the material covered in the book.

Information theory has proved to be effective for solving many computer vision and pattern recognition (CVPR) problems (such as image matching, clustering and segmentation, saliency detection, feature selection, optimal classifier design and many others). Nowadays, researchers are widely bringing information theory elements to the CVPR arena.

Among these elements there are measures (entropy, mutual information...), principles (maximum entropy, minimax entropy...) and theories (rate distortion theory, method of types...). This book explores and introduces the latter elements through an incremental complexity approach at the same time where CVPR problems are formulated and the most representative algorithms are presented. Interesting connections between information theory principles when applied to different problems are highlighted, seeking a comprehensive research roadmap. The result is a novel tool both for CVPR and machine learning researchers, and contributes to a cross-fertilization of both areas.

This textbook provides an accessible general introduction to the essential topics in computer vision. Classroom-tested programming exercises and review questions are also supplied at the end of each chapter. Features: provides an introduction to the basic notation and mathematical concepts for describing an image and the key concepts for mapping an image into an image; explains the topologic and geometric basics for analysing image regions and distributions of image values and discusses identifying patterns in an image; introduces optic flow for representing dense motion and various topics in sparse motion analysis; describes special approaches for image binarization and segmentation of still images or video frames; examines the basic components of a computer vision system; reviews different techniques for vision-based 3D shape reconstruction; includes a discussion of stereo matchers and the phase-congruency model for image features; presents an introduction into classification and learning.

Learn how to build your own computer vision (CV) applications quickly and easily with SimpleCV, an open source framework written in Python. Through examples of real-world applications, this hands-on guide introduces you to basic CV techniques for collecting, processing, and analyzing streaming digital images. You'll then learn how to apply these methods with SimpleCV, using sample Python code. All you need to get started is a Windows, Mac, or Linux system, and a willingness to put CV to work in a variety of ways.

Programming experience is optional. Capture images from several sources, including webcams, smartphones, and Kinect Filter image input so your application processes only necessary information Manipulate images by performing basic arithmetic on pixel values Use feature detection techniques to focus on interesting parts of an image Work with several features in a single image, using the NumPy and SciPy Python libraries Learn about optical flow to identify objects that change between two image frames Use SimpleCV's command line and code editor to run examples and test techniques

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.

Image algebra is a comprehensive, unifying theory of image transformations, image analysis, and image understanding. In 1996, the bestselling first edition of the Handbook of Computer Vision Algorithms in Image Algebra introduced engineers, scientists, and students to this powerful tool, its basic concepts, and its use in the concise representation of computer vision algorithms. Updated to reflect recent developments and advances, the second edition continues to provide an outstanding introduction to image algebra. It describes more than 80 fundamental computer vision techniques and introduces the portable iaC++ library, which supports image algebra programming in the C++ language. Revisions to the first edition include a new chapter on geometric manipulation and spatial transformation, several additional algorithms, and the addition of exercises to each chapter. The authors-both instrumental in the groundbreaking development of image algebra-introduce each technique with a brief discussion of its purpose and methodology, then provide its precise mathematical formulation. In addition to furnishing the simple yet powerful utility of image algebra, the Handbook of Computer Vision Algorithms in Image Algebra supplies the core of knowledge all computer vision practitioners need. It offers a more practical, less esoteric presentation than those found in research publications that will soon earn it a prime location on your reference shelf.

This comprehensive guide provides a uniquely practical, application-focused introduction to medical image analysis. This fully updated new edition has been enhanced with material on the latest developments in the field, whilst retaining the original focus on segmentation, classification and registration. Topics and features: presents learning objectives, exercises and concluding remarks in each chapter; describes a range of common imaging techniques, reconstruction techniques and image artifacts, and discusses the archival and transfer of images; reviews an expanded selection of techniques for image enhancement, feature detection, feature generation, segmentation, registration, and validation; examines analysis methods in view of image-based guidance in the operating room (NEW); discusses the use of deep convolutional networks for segmentation and labeling tasks (NEW); includes appendices on Markov random field optimization, variational calculus and principal component analysis.

Explains the theory behind basic computer vision and provides a bridge from the theory to practical implementation using the industry standard OpenCV libraries Computer Vision is a rapidly expanding area and it is becoming progressively easier for developers to make use of this field due to the ready availability of high quality libraries (such as OpenCV2). This text is intended to facilitate the practical use of computer vision with the goal being to bridge the gap between the theory and the practical implementation of computer vision.

The book will explain how to use the relevant OpenCV library routines and will be accompanied by a full working program including the code snippets from the text. This textbook is a heavily illustrated, practical introduction to an exciting field, the applications of which are becoming almost ubiquitous. We are now surrounded by cameras, for example cameras on computers & tablets/ cameras built into our mobile phones/ cameras in games consoles; cameras imaging difficult modalities (such as ultrasound, X-ray, MRI) in hospitals, and surveillance cameras. This book is concerned with helping the next generation of computer developers to make use of all these images in order to develop systems which are more intuitive and interact with us in more intelligent ways. Explains the theory behind basic computer vision and provides a bridge from the theory to practical implementation using the industry standard OpenCV libraries. Offers an introduction to computer vision, with enough theory to make clear how the various algorithms work but with an emphasis on practical programming issues. Provides enough material for a one semester course in computer vision at senior undergraduate and Masters levels. Includes the basics of cameras and images and image processing to remove noise, before moving on to topics such as image histogramming; binary imaging; video processing to detect and model moving objects; geometric operations & camera models; edge detection; features detection; recognition in images. Contains a large number of vision application problems to provide students with the opportunity to solve real problems. Images or videos for these problems are provided in the resources associated with this book which include an enhanced eBook.

In 1936, when he was just twenty-four years old, Alan Turing wrote a remarkable paper in which he outlined the theory of computation, laying out the ideas that underlie all modern computers. This groundbreaking and powerful theory now forms the basis of computer science. In Turing's Vision, Chris Bernhardt explains the theory, Turing's most important contribution, for the general reader. Bernhardt argues that the strength of Turing's theory is its simplicity, and that, explained in a straightforward manner, it is eminently understandable by the nonspecialist. As Marvin Minsky writes, "The sheer simplicity of the theory's foundation and extraordinary short path from this foundation to its logical and surprising conclusions give the theory a mathematical beauty that alone guarantees it a permanent place in computer theory." Bernhardt begins with the foundation and systematically builds to the surprising conclusions. He also views Turing's theory in the context of mathematical history, other views of computation (including those of Alonzo Church), Turing's later work, and the birth of the modern computer. In the paper, "On Computable Numbers, with an Application to the Entscheidungsproblem," Turing thinks carefully about how humans perform computation, breaking it down into a sequence of steps, and then constructs theoretical machines capable of performing each step. Turing wanted to show that there were problems that were beyond any computer's ability to solve; in particular, he wanted to find a decision problem that he could prove was undecidable. To explain Turing's ideas, Bernhardt examines three well-known decision problems to explore the concept of undecidability; investigates theoretical computing machines, including Turing machines; explains universal machines; and proves that certain problems are undecidable, including Turing's problem concerning computable numbers.

Concise Computer Vision: An Introduction into Theory and Algorithms Springer

Annotation. Computer and Machine Vision: Theory, Algorithms, Practicalities (previously entitled Machine Vision) clearly and systematically presents the basic methodology of computer and machine vision, covering the essential elements of the theory while emphasizing algorithmic and practical design constraints. This fully revised fourth edition has brought in more of the concepts and applications of computer vision, making it a very comprehensive and up-to-date tutorial text suitable for graduate students, researchers and R the first of these has been widely used internationally for more than 20 years, and is now out in this much enhanced fourth edition. Roy holds a DSc at the University of London, and has been awarded Distinguished Fellow of the British Machine Vision Association, and Fellow of the International Association of Pattern Recognition. Mathematics and essential theory are made approachable by careful explanations and well-illustrated examples. Updated content and new sections cover topics such as human iris location, image stitching, line detection using RANSAC, performance measures, and hyperspectral imaging. The 'recent developments' section now included in each chapter will be useful in bringing students and practitioners up to date with the subject.

This book explains the theory and application of evolutionary computer vision, a new paradigm where challenging vision problems can be approached using the techniques of evolutionary computing. This methodology achieves excellent results for defining fitness functions and representations for problems by merging evolutionary computation with mathematical optimization to produce automatic creation of emerging visual behaviors. In the first part of the book the author surveys the literature in concise form, defines the relevant terminology, and offers historical and philosophical motivations for the key research problems in the field. For researchers from the computer vision community, he offers a simple introduction to the evolutionary computing paradigm. The second part of the book focuses on implementing evolutionary algorithms that solve given problems using working programs in the major fields of low-, intermediate- and high-level computer vision. This book will be of value to researchers, engineers, and students in the fields of computer vision, evolutionary computing, robotics, biologically inspired mechatronics, electronics engineering, control, and artificial intelligence.

A cookbook of algorithms for common image processing applications. Thanks to advances in computer hardware and software, algorithms have been developed that support sophisticated image processing without requiring an extensive background in mathematics. This bestselling book has been fully updated with the newest of these, including 2D vision methods in content-based searches and the use of graphics cards as image processing computational aids. It's an ideal reference for software engineers and developers, advanced programmers, graphics programmers, scientists, and other specialists who require highly specialized image processing. Algorithms now exist for a wide variety of sophisticated image processing applications required by software engineers and developers, advanced programmers, graphics programmers, scientists, and related specialists. This bestselling book has been completely updated to include the latest algorithms, including 2D vision methods in content-based searches, details on modern classifier methods, and graphics cards used as image processing computational aids. Saves hours of mathematical calculating by using distributed processing and GPU programming, and gives non-mathematicians the shortcuts needed to program relatively sophisticated applications. Algorithms for Image Processing and Computer Vision, 2nd Edition provides the tools to speed development of image processing applications.

Discover how CUDA allows OpenCV to handle complex and rapidly growing image data processing in computer and machine vision by accessing the power of GPU. Key Features: Explore examples to leverage the GPU processing power with OpenCV and CUDA. Enhance the performance of algorithms on embedded hardware platforms. Discover C++ and Python libraries for GPU acceleration. Book Description: Computer vision has been revolutionizing a wide range of industries, and OpenCV is the most widely chosen tool for computer vision with its ability to work in multiple programming languages. Nowadays, in computer vision, there is a need to process large images in real time, which is difficult to handle for OpenCV on its own. This is where CUDA comes into the picture, allowing OpenCV to leverage powerful NVIDIA GPUs. This book provides a detailed overview of integrating OpenCV with CUDA for practical applications. To start with, you'll understand GPU programming with CUDA, an

essential aspect for computer vision developers who have never worked with GPUs. You'll then move on to exploring OpenCV acceleration with GPUs and CUDA by walking through some practical examples. Once you have got to grips with the core concepts, you'll familiarize yourself with deploying OpenCV applications on NVIDIA Jetson TX1, which is popular for computer vision and deep learning applications. The last chapters of the book explain PyCUDA, a Python library that leverages the power of CUDA and GPUs for accelerations and can be used by computer vision developers who use OpenCV with Python. By the end of this book, you'll have enhanced computer vision applications with the help of this book's hands-on approach. What you will learn Understand how to access GPU device properties and capabilities from CUDA programs Learn how to accelerate searching and sorting algorithms Detect shapes such as lines and circles in images Explore object tracking and detection with algorithms Process videos using different video analysis techniques in Jetson TX1 Access GPU device properties from the PyCUDA program Understand how kernel execution works Who this book is for This book is a go-to guide for you if you are a developer working with OpenCV and want to learn how to process more complex image data by exploiting GPU processing. A thorough understanding of computer vision concepts and programming languages such as C++ or Python is expected.

This book is an essential tool for second-year undergraduate students and above, providing clear and concise explanations of the basic concepts of computer graphics, and enabling the reader to immediately implement these concepts in Java 2D and/or 3D with only elementary knowledge of the programming language. Features: provides an ideal, self-contained introduction to computer graphics, with theory and practice presented in integrated combination; presents a practical guide to basic computer graphics programming using Java 2D and 3D; includes new and expanded content on the integration of text in 3D, particle systems, billboard behaviours, dynamic surfaces, the concept of level of detail, and the use of functions of two variables for surface modelling; contains many pedagogical tools, including numerous easy-to-understand example programs and end-of-chapter exercises; supplies useful supplementary material, including additional exercises, solutions, and program examples, at an associated website.

This book covers the fundamentals in designing and deploying techniques using deep architectures. It is intended to serve as a beginner's guide to engineers or students who want to have a quick start on learning and/or building deep learning systems. This book provides a good theoretical and practical understanding and a complete toolkit of basic information and knowledge required to understand and build convolutional neural networks (CNN) from scratch. The book focuses explicitly on convolutional neural networks, filtering out other material that co-occur in many deep learning books on CNN topics.

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

A basic problem in computer vision is to understand the structure of a real world scene given several images of it. Techniques for solving this problem are taken from projective geometry and photogrammetry. Here, the authors cover the geometric principles and their algebraic representation in terms of camera projection matrices, the fundamental matrix and the trifocal tensor. The theory and methods of computation of these entities are discussed with real examples, as is their use in the reconstruction of scenes from multiple images. The new edition features an extended introduction covering the key ideas in the book (which itself has been updated with additional examples and appendices) and significant new results which have appeared since the first edition. Comprehensive background material is provided, so readers familiar with linear algebra and basic numerical methods can understand the projective geometry and estimation algorithms presented, and implement the algorithms directly from the book.

A cookbook of the hottest new algorithms and cutting-edge techniques in image processing and computer vision This amazing book/CD package puts the power of all the hottest new image processing techniques and algorithms in your hands. Based on J. R. Parker's exhaustive survey of Internet newsgroups worldwide, Algorithms for Image Processing and Computer Vision answers the most frequently asked questions with practical solutions. Parker uses dozens of real-life examples taken from fields such as robotics, space exploration, forensic analysis, cartography, and medical diagnostics, to clearly describe the latest techniques for morphing, advanced edge detection, wavelets, texture classification, image restoration, symbol recognition, and genetic algorithms, to name just a few. And, best of all, he implements each method covered in C and provides all the source code on the CD. For the first time, you're rescued from the hours of mind-numbing mathematical calculations it would ordinarily take to program these state-of-the-art image processing capabilities into software. At last, nonmathematicians get all the shortcuts they need for sophisticated image recognition and processing applications. On the CD-ROM you'll find: * Complete code for examples in the book * A gallery of images illustrating the results of advanced techniques * A free GNU compiler that lets you run source code on any platform * A system for restoring damaged or blurred images * A genetic algorithms package

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

This text provides readers with a starting point to understand and investigate the literature of computer vision, listing conferences, journals and Internet sites.

Directed at third-year undergraduates and postgraduates this challenging text offers a complete and up-to-date introduction to computer vision and image processing. Examples in both Java and C++ are used throughout the book making it suitable for a wide range of courses. Topics covered include: image representation, feature extraction, feature recognition, architecture of computer vision systems, tracking moving objects, and image coding.

Feature Extraction for Image Processing and Computer Vision is an essential guide to the implementation of image processing and computer vision techniques, with tutorial introductions and sample code in MATLAB and Python. Algorithms are presented and fully explained to enable complete understanding of the methods and techniques demonstrated. As one reviewer noted, "The main strength of the proposed book is the link between theory and exemplar code of the algorithms." Essential background theory is carefully explained. This text gives students and researchers in image processing and computer vision a complete introduction to classic and state-of-the-art methods in feature extraction together with practical guidance on their implementation. The only text to concentrate on feature extraction with working implementation and worked through mathematical derivations and algorithmic methods A thorough overview of

available feature extraction methods including essential background theory, shape methods, texture and deep learning Up to date coverage of interest point detection, feature extraction and description and image representation (including frequency domain and colour) Good balance between providing a mathematical background and practical implementation Detailed and explanatory of algorithms in MATLAB and Python

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

[Copyright: e4726f8042a67a33f0b5202e2cda37ec](#)