

3d Deep Shape Descriptor Cv Foundation

This textbook is designed for postgraduate studies in the field of 3D Computer Vision. It also provides a useful reference for industrial practitioners; for example, in the areas of 3D data capture, computer-aided geometric modelling and industrial quality assurance. This second edition is a significant upgrade of existing topics with novel findings. Additionally, it has new material covering consumer-grade RGB-D cameras, 3D morphable models, deep learning on 3D datasets, as well as new applications in the 3D digitization of cultural heritage and the 3D phenotyping of crops. Overall, the book covers three main areas: ? 3D imaging, including passive 3D imaging, active triangulation 3D imaging, active time-of-flight 3D imaging, consumer RGB-D cameras, and 3D data representation and visualisation; ? 3D shape analysis, including local descriptors, registration, matching, 3D morphable models, and deep learning on 3D datasets; and ? 3D applications, including 3D face recognition, cultural heritage and 3D phenotyping of plants. 3D computer vision is a rapidly advancing area in computer science. There are many real-world applications that demand high-performance 3D imaging and analysis and, as a result, many new techniques and commercial products have been developed. However, many challenges remain on how to analyse the captured data in a way that is sufficiently fast, robust and accurate for the application. Such challenges include metrology, semantic segmentation, classification and recognition. Thus, 3D imaging, analysis and their applications remain a highly-active research field that will continue to attract intensive attention from the research community with the ultimate goal of fully automating the 3D data capture, analysis and inference pipeline.

This book presents best selected research papers presented at the Thirteenth International Conference on Applied Mathematics and Mechanics in the Aerospace Industry (AMMAI 2020), held from September 6 to September 13, 2020, at the Alushta Health and Educational Center (The Republic of Crimea). The book is dedicated to solving actual problems of applied mechanics using modern computer technology including smart paradigms. Physical and mathematical models, numerical methods, computational algorithms, and software complexes are discussed, which allow to carry out high-precision mathematical modeling in fluid, gas, and plasma mechanics, in general mechanics, deformable solid mechanics, in strength, destruction and safety of structures, etc. Technologies and software systems that provide effective solutions to the problems at various multi-scale levels are considered. Special attention is paid to the training of highly qualified specialists for the aviation and space industry. The book is recommended for specialists in the field of applied mathematics and mechanics, mathematical modeling, information technologies, and developers of modern applied software systems.

This book constitutes the proceedings of the 6th International Workshop on Computational Topology in Image Context, CTIC 2016, held in Marseille, France, in June 2016. The 24 papers presented in this volume were carefully reviewed and selected from 35 submissions. Additionally, this volume contains 2 invited papers. CTIC covers a wide range of topics such as: topological invariants and their computation, homology, cohomology, linking number, fundamental groups; algorithm optimization in discrete geometry, transfer of mathematical tools, parallel computation in multi-dimensional volume context, hierarchical approaches; experimental evaluation of algorithms and heuristics; combinatorial or multi-resolution models; discrete or computational topology; geometric modeling guided by topological constraints; computational topological dynamics; and use of topological information in discrete geometry applications.

This book discusses recent advances in object detection and recognition using deep learning methods, which have achieved great success in the field of computer vision and image processing. It provides a systematic and methodical overview of the latest developments in deep learning theory and its applications to computer vision, illustrating them using key topics, including object detection, face analysis, 3D object recognition, and image retrieval. The book offers a rich blend of theory and practice. It is suitable for students, researchers and practitioners interested in deep learning, computer vision and beyond and can also be used as a reference book. The comprehensive comparison of various deep-learning applications helps readers with a basic understanding of machine learning and calculus grasp the theories and inspires applications in other computer vision tasks.

An in-depth description of the state-of-the-art of 3D shape analysis techniques and their applications This book discusses the different topics that come under the title of "3D shape analysis". It covers the theoretical foundations and the major solutions that have been presented in the literature. It also establishes links between solutions proposed by different communities that studied 3D shape, such as mathematics and statistics, medical imaging, computer vision, and computer graphics. The first part of 3D Shape Analysis: Fundamentals, Theory, and Applications provides a review of the background concepts such as methods for the acquisition and representation of 3D geometries, and the fundamentals of geometry and topology. It specifically covers stereo matching, structured light, and intrinsic vs. extrinsic properties of shape. Parts 2 and 3 present a range of mathematical and algorithmic tools (which are used for e.g., global descriptors, keypoint detectors, local feature descriptors, and algorithms) that are commonly used for the detection, registration, recognition, classification, and retrieval of 3D objects. Both also place strong emphasis on recent techniques motivated by the spread of commodity devices for 3D acquisition. Part 4 demonstrates the use of these techniques in a selection of 3D shape analysis applications. It covers 3D face recognition, object recognition in 3D scenes, and 3D shape retrieval. It also discusses examples of semantic applications and cross domain 3D retrieval, i.e. how to retrieve 3D models using various types of modalities, e.g. sketches and/or images. The book concludes with a summary of the main ideas and discussions of the future trends. 3D Shape Analysis: Fundamentals, Theory, and Applications is an excellent reference for graduate students, researchers, and professionals in different fields of mathematics, computer science, and engineering. It is also ideal for courses in computer vision and computer graphics, as well as for those seeking 3D industrial/commercial solutions.

This book constitutes the proceedings of the Third International Workshop on Similarity Based Pattern Analysis and Recognition, SIMBAD 2015, which was held in Copenhagen, Denmark, in October 2015. The 15 full and 8 short papers presented were carefully reviewed and selected from 30 submissions. The workshop focus on problems, techniques, applications, and perspectives: from supervised to unsupervised learning, from generative to discriminative models, and from theoretical issues to empirical validations.

Image registration is the process of systematically placing separate images in a common frame of reference so that the information they contain can be optimally integrated or compared. This is becoming the central tool for image analysis, understanding, and visualization in both medical and scientific applications. Medical Image Registration provid

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including

artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

This book constitutes the refereed proceedings of CVM 2012, the First International Conference on Computational Visual Media, held in Beijing, China, in November 2012. The 33 revised full papers were carefully reviewed and selected from 81 submissions. The papers are organized in topical sections on image processing I and II, geometric processing, saliency, recognition, perception and learning, shape analysis, media retrieval, and capture, rendering and visualization.

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.

Computer Vision – ACCV 2018 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V Springer

Discover interesting recipes to help you understand the concepts of object detection, image processing, and facial detection
Key Features Explore the latest features and APIs in OpenCV 4 and build computer vision algorithms
Develop effective, robust, and fail-safe vision for your applications
Build computer vision algorithms with machine learning capabilities
Book Description OpenCV is an image and video processing library used for all types of image and video analysis. Throughout the book, you'll work through recipes that implement a variety of tasks, such as facial recognition and detection. With 70 self-contained tutorials, this book examines common pain points and best practices for computer vision (CV) developers. Each recipe addresses a specific problem and offers a proven, best-practice solution with insights into how it works, so that you can copy the code and configuration files and modify them to suit your needs. This book begins by setting up OpenCV, and explains how to manipulate pixels. You'll understand how you can process images with classes and count pixels with histograms. You'll also learn detecting, describing, and matching interest points. As you advance through the chapters, you'll get to grips with estimating projective relations in images, reconstructing 3D scenes, processing video sequences, and tracking visual motion. In the final chapters, you'll cover deep learning concepts such as face and object detection. By the end of the book, you'll be able to confidently implement a range of computer vision algorithms to meet the technical requirements of your complex CV projects
What you will learn
Install and create a program using the OpenCV library
Segment images into homogenous regions and extract meaningful objects
Apply image filters to enhance image content
Exploit image geometry to relay different views of a pictured scene
Calibrate the camera from different image observations
Detect people and objects in images using machine learning techniques
Reconstruct a 3D scene from images
Explore face detection using deep learning
Who this book is for
If you're a CV developer or professional who already uses or would like to use OpenCV for building computer vision software, this book is for you. You'll also find this book useful if you're a C++ programmer looking to extend your computer vision skillset by learning OpenCV.

Theoretical results suggest that in order to learn the kind of complicated functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

This book describes new theories and applications of artificial neural networks, with a special focus on answering questions in neuroscience, biology and biophysics and cognitive

research. It covers a wide range of methods and technologies, including deep neural networks, large scale neural models, brain computer interface, signal processing methods, as well as models of perception, studies on emotion recognition, self-organization and many more. The book includes both selected and invited papers presented at the XXI International Conference on Neuroinformatics, held on October 7-11, 2019, in Dolgoprudny, a town in Moscow region, Russia.

One of the grand challenges of artificial intelligence is to enable computers to interpret 3D scenes and objects from imagery. This book organizes and introduces major concepts in 3D scene and object representation and inference from still images, with a focus on recent efforts to fuse models of geometry and perspective with statistical machine learning. The book is organized into three sections: (1) Interpretation of Physical Space; (2) Recognition of 3D Objects; and (3) Integrated 3D Scene Interpretation. The first discusses representations of spatial layout and techniques to interpret physical scenes from images. The second section introduces representations for 3D object categories that account for the intrinsically 3D nature of objects and provide robustness to change in viewpoints. The third section discusses strategies to unite inference of scene geometry and object pose and identity into a coherent scene interpretation. Each section broadly surveys important ideas from cognitive science and artificial intelligence research, organizes and discusses key concepts and techniques from recent work in computer vision, and describes a few sample approaches in detail. Newcomers to computer vision will benefit from introductions to basic concepts, such as single-view geometry and image classification, while experts and novices alike may find inspiration from the book's organization and discussion of the most recent ideas in 3D scene understanding and 3D object recognition. Specific topics include: mathematics of perspective geometry; visual elements of the physical scene, structural 3D scene representations; techniques and features for image and region categorization; historical perspective, computational models, and datasets and machine learning techniques for 3D object recognition; inferences of geometrical attributes of objects, such as size and pose; and probabilistic and feature-passing approaches for contextual reasoning about 3D objects and scenes. Table of Contents: Background on 3D Scene Models / Single-view Geometry / Modeling the Physical Scene / Categorizing Images and Regions / Examples of 3D Scene Interpretation / Background on 3D Recognition / Modeling 3D Objects / Recognizing and Understanding 3D Objects / Examples of 2D 1/2 Layout Models / Reasoning about Objects and Scenes / Cascades of Classifiers / Conclusion and Future Directions

It is our belief that researchers and practitioners acquire, through experience and word-of-mouth, techniques and heuristics that help them successfully apply neural networks to difficult real world problems. Often these "tricks" are theoretically well motivated. Sometimes they are the result of trial and error. However, their most common link is that they are usually hidden in people's heads or in the back pages of space-constrained conference papers. As a result newcomers to the field waste much time wondering why their networks train so slowly and perform so poorly. This book is an outgrowth of a 1996 NIPS workshop called Tricks of the Trade whose goal was to begin the process of gathering and documenting these tricks. The interest that the workshop generated motivated us to expand our collection and compile it into this book. Although we have no doubt that there are many tricks we have missed, we hope that what we have included will prove to be useful, particularly to those who are relatively new to the field. Each chapter contains one or more tricks presented by a given author (or authors). We have attempted to group related chapters into sections, though we recognize that the different sections are far from disjoint. Some of the chapters (e.g., 1, 13, 17) contain entire systems of tricks that are far more general than the category they have been placed in.

This book is a guide for researchers and practitioners to the new frontiers of 3D shape analysis and the complex mathematical tools most methods rely on. The target reader includes students, researchers and professionals with an undergraduate mathematics background, who wish to understand the mathematics behind shape analysis. The authors begin with a quick review of basic concepts in geometry, topology, differential geometry, and proceed to advanced notions of algebraic topology, always keeping an eye on the application of the theory, through examples of shape analysis methods such as 3D segmentation, correspondence, and retrieval. A number of research solutions in the field come from advances in pure and applied mathematics, as well as from the re-reading of classical theories and their adaptation to the discrete setting. In a world where disciplines (fortunately) have blurred boundaries, the authors believe that this guide will help to bridge the distance between theory and practice. Table of Contents: Acknowledgments / Figure Credits / About this Book / 3D Shape Analysis in a Nutshell / Geometry, Topology, and Shape Representation / Differential Geometry and Shape Analysis / Spectral Methods for Shape Analysis / Maps and Distances between Spaces / Algebraic Topology and Topology Invariants / Differential Topology and Shape Analysis / Reeb Graphs / Morse and Morse-Smale Complexes / Topological Persistence / Beyond Geometry and Topology / Resources / Bibliography / Authors' Biographies

The three-volume set, consisting of LNCS 9008, 9009, and 9010, contains carefully reviewed and selected papers presented at 15 workshops held in conjunction with the 12th Asian Conference on Computer Vision, ACCV 2014, in Singapore, in November 2014. The 153 full papers presented were selected from numerous submissions. LNCS 9008 contains the papers selected for the Workshop on Human Gait and Action Analysis in the Wild, the Second International Workshop on Big Data in 3D Computer Vision, the Workshop on Deep Learning on Visual Data, the Workshop on Scene Understanding for Autonomous Systems, and the Workshop on Robust Local Descriptors for Computer Vision. LNCS 9009 contains the papers selected for the Workshop on Emerging Topics on Image Restoration and Enhancement, the First International Workshop on Robust Reading, the Second Workshop on User-Centred Computer Vision, the International Workshop on Video Segmentation in Computer Vision, the Workshop: My Car Has Eyes: Intelligent Vehicle with Vision Technology, the Third Workshop on E-Heritage, and the Workshop on Computer Vision for Affective Computing. LNCS 9010 contains the papers selected for the Workshop on Feature and Similarity for Computer Vision, the Third International Workshop on Intelligent Mobile and Egocentric Vision, and the Workshop on Human Identification for Surveillance.

The first book of its kind dedicated to the challenge of person re-identification, this text provides an in-depth, multidisciplinary discussion of recent developments and state-of-the-art methods. Features: introduces examples of robust feature representations, reviews salient feature weighting and selection mechanisms and examines the benefits of semantic attributes; describes how

to segregate meaningful body parts from background clutter; examines the use of 3D depth images and contextual constraints derived from the visual appearance of a group; reviews approaches to feature transfer function and distance metric learning and discusses potential solutions to issues of data scalability and identity inference; investigates the limitations of existing benchmark datasets, presents strategies for camera topology inference and describes techniques for improving post-rank search efficiency; explores the design rationale and implementation considerations of building a practical re-identification system.

An evolving, living organic/inorganic covering, soil is in dynamic equilibrium with the atmosphere above, the biosphere within, and the geology below. It acts as an anchor for roots, a purveyor of water and nutrients, a residence for a vast community of microorganisms and animals, a sanitizer of the environment, and a source of raw materials for co

World-renowned economist Klaus Schwab, Founder and Executive Chairman of the World Economic Forum, explains that we have an opportunity to shape the fourth industrial revolution, which will fundamentally alter how we live and work. Schwab argues that this revolution is different in scale, scope and complexity from any that have come before. Characterized by a range of new technologies that are fusing the physical, digital and biological worlds, the developments are affecting all disciplines, economies, industries and governments, and even challenging ideas about what it means to be human. Artificial intelligence is already all around us, from supercomputers, drones and virtual assistants to 3D printing, DNA sequencing, smart thermostats, wearable sensors and microchips smaller than a grain of sand. But this is just the beginning: nanomaterials 200 times stronger than steel and a million times thinner than a strand of hair and the first transplant of a 3D printed liver are already in development. Imagine "smart factories" in which global systems of manufacturing are coordinated virtually, or implantable mobile phones made of biosynthetic materials. The fourth industrial revolution, says Schwab, is more significant, and its ramifications more profound, than in any prior period of human history. He outlines the key technologies driving this revolution and discusses the major impacts expected on government, business, civil society and individuals. Schwab also offers bold ideas on how to harness these changes and shape a better future--one in which technology empowers people rather than replaces them; progress serves society rather than disrupts it; and in which innovators respect moral and ethical boundaries rather than cross them. We all have the opportunity to contribute to developing new frameworks that advance progress.

Computer Vision Metrics provides an extensive survey and analysis of over 100 current and historical feature description and machine vision methods, with a detailed taxonomy for local, regional and global features. This book provides necessary background to develop intuition about why interest point detectors and feature descriptors actually work, how they are designed, with observations about tuning the methods for achieving robustness and invariance targets for specific applications. The survey is broader than it is deep, with over 540 references provided to dig deeper. The taxonomy includes search methods, spectra components, descriptor representation, shape, distance functions, accuracy, efficiency, robustness and invariance attributes, and more. Rather than providing 'how-to' source code examples and shortcuts, this book provides a counterpoint discussion to the many fine opencv community source code resources available for hands-on practitioners.

This book seeks to bridge the gap between statistics and computer science. It provides an overview of Monte Carlo methods, including Sequential Monte Carlo, Markov Chain Monte Carlo, Metropolis-Hastings, Gibbs Sampler, Cluster Sampling, Data Driven MCMC, Stochastic Gradient descent, Langevin Monte Carlo, Hamiltonian Monte Carlo, and energy landscape mapping. Due to its comprehensive nature, the book is suitable for developing and teaching graduate courses on Monte Carlo methods. To facilitate learning, each chapter includes several representative application examples from various fields. The book pursues two main goals: (1) It introduces researchers to applying Monte Carlo methods to broader problems in areas such as Computer Vision, Computer Graphics, Machine Learning, Robotics, Artificial Intelligence, etc.; and (2) it makes it easier for scientists and engineers working in these areas to employ Monte Carlo methods to enhance their research.

This book constitutes the refereed proceedings of the 19th Scandinavian Conference on Image Analysis, SCIA 2015, held in Copenhagen, Denmark, in June 2015. The 45 revised papers presented were carefully reviewed and selected from 67 submissions. The contributions are structured in topical sections on novel applications of vision systems, pattern recognition, machine learning, feature extraction, segmentation, 3D vision to medical and biomedical image analysis.

Each chapter in the book is an individual project and each project is constructed with step-by-step instructions, clearly explained code, and includes the necessary screenshots. You should have basic OpenCV and C/C++ programming experience before reading this book, as it is aimed at Computer Science graduates, researchers, and computer vision experts widening their expertise.

The 30-volume set, comprising the LNCS books 12346 until 12375, constitutes the refereed proceedings of the 16th European Conference on Computer Vision, ECCV 2020, which was planned to be held in Glasgow, UK, during August 23-28, 2020. The conference was held virtually due to the COVID-19 pandemic. The 1360 revised papers presented in these proceedings were carefully reviewed and selected from a total of 5025 submissions. The papers deal with topics such as computer vision; machine learning; deep neural networks; reinforcement learning; object recognition; image classification; image processing; object detection; semantic segmentation; human pose estimation; 3d reconstruction; stereo vision; computational photography; neural networks; image coding; image reconstruction; object recognition; motion estimation.

The six volume set LNCS 11361-11366 constitutes the proceedings of the 14th Asian Conference on Computer Vision, ACCV 2018, held in Perth, Australia, in December 2018. The total of 274 contributions was carefully reviewed and selected from 979 submissions during two rounds of reviewing and improvement. The papers focus on motion and tracking, segmentation and grouping, image-based modeling, deep learning, object recognition object recognition, object detection and categorization, vision and language, video analysis and event recognition, face and gesture analysis, statistical methods and learning, performance evaluation, medical image analysis, document analysis, optimization methods, RGBD and depth camera processing, robotic vision, applications of computer vision.

Graphs are useful data structures in complex real-life applications such as modeling physical systems, learning molecular fingerprints, controlling traffic networks, and recommending friends in social networks. However, these tasks require dealing with non-Euclidean graph data that contains rich relational information between elements and cannot be well handled by traditional deep learning models (e.g., convolutional neural networks (CNNs) or recurrent neural networks (RNNs)). Nodes in graphs usually contain useful feature information that cannot be well addressed in most unsupervised representation learning methods (e.g., network embedding methods). Graph neural networks (GNNs) are proposed to combine the feature information and the graph structure to learn better representations on graphs via feature propagation and aggregation. Due to its convincing performance and high interpretability, GNN has recently become a widely applied graph analysis tool. This book provides a comprehensive

introduction to the basic concepts, models, and applications of graph neural networks. It starts with the introduction of the vanilla GNN model. Then several variants of the vanilla model are introduced such as graph convolutional networks, graph recurrent networks, graph attention networks, graph residual networks, and several general frameworks. Variants for different graph types and advanced training methods are also included. As for the applications of GNNs, the book categorizes them into structural, non-structural, and other scenarios, and then it introduces several typical models on solving these tasks. Finally, the closing chapters provide GNN open resources and the outlook of several future directions.

"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.

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.

The LNCS volume 11818 constitutes the proceedings of the 14th Chinese Conference on Biometric Recognition, held in Zhuzhou, China, in October 2019. The 56 papers presented in this book were carefully reviewed and selected from 74 submissions. The papers cover a wide range of topics such as face recognition and analysis; hand-based biometrics; eye-based biometrics; gesture, gait, and action; emerging biometrics; feature extraction and classification theory; and behavioral biometrics.

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

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

"John Russ is the master of explaining how image processing gets applied to real-world situations. With Brent Neal, he's done it again in Measuring Shape, this time explaining an expanded toolbox of techniques that includes useful, state-of-the-art methods that can be applied to the broad problem of understanding, characterizing, and measuring shape. He has a gift for finding the kernel of a particular algorithm, explaining it in simple terms, then giving concrete examples that are easily understood. His perspective comes from solving real-world problems and separating out what works in practice from what is just an abstract curiosity." —Tom Malzbender, Hewlett-Packard Laboratories, Palo Alto, California, USA Useful for those working in fields including industrial quality control, research, and security applications, Measuring Shape is a handbook for the practical application of shape measurement. Covering a wide range of shape measurements likely to be encountered in the literature and in software packages, this book presents an intentionally diverse set of examples that illustrate and enable readers to compare methods used for measurement and quantitative description of 2D and 3D shapes. It stands apart through its focus on examples and applications, which help readers quickly grasp the usefulness of presented techniques without having to approach them through the underlying mathematics. An elusive concept, shape is a principal governing factor in determining the behavior of objects and structures. Essential to recognizing and classifying objects, it is the central link in manmade and natural processes. Shape dictates everything from the stiffness of a construction beam, to the ability of a leaf to catch water, to the marketing and packaging of consumer products. This book emphasizes techniques that are quantitative and produce a meaningful yet compact set of numerical values that can be used for statistical analysis, comparison, correlation, classification, and identification. Written by two renowned authors from both industry and academia, this resource explains why users should select a particular method, rather than simply discussing how to use it. Showcasing each process in a clear, accessible, and well-organized way, they explore why a particular one might be appropriate in a given situation, yet a poor choice in another. Providing extensive examples, plus full mathematical descriptions of the various measurements involved, they detail the advantages and limitations of each method and explain the ways they can be implemented to discover important correlations between shape and object history or behavior. This uncommon assembly of information also includes sets of data on real-world objects that are used to compare the performance and utility of the various presented approaches.

The two volumes LNCS 11935 and 11936 constitute the proceedings of the 9th International Conference on Intelligence Science and Big Data Engineering, IScIDE 2019, held in Nanjing, China, in October 2019. The 84 full papers presented were carefully reviewed and selected from 252 submissions. The papers are organized in two parts: visual data engineering; and big data and machine learning. They cover a large range of topics including information theoretic and Bayesian approaches, probabilistic graphical models, big data analysis, neural networks and neuro-informatics, bioinformatics, computational biology and brain-computer interfaces, as well as advances in fundamental pattern recognition techniques relevant to image processing, computer vision and machine learning.

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Where To Download 3d Deep Shape Descriptor Cv Foundation

The three-volume set, consisting of LNCS 9008, 9009, and 9010, contains carefully reviewed and selected papers presented at 15 workshops held in conjunction with the 12th Asian Conference on Computer Vision, ACCV 2014, in Singapore, in November 2014. The 153 full papers presented were selected from numerous submissions. LNCS 9008 contains the papers selected for the Workshop on Human Gait and Action Analysis in the Wild, the Second International Workshop on Big Data in 3D Computer Vision, the Workshop on Deep Learning on Visual Data, the Workshop on Scene Understanding for Autonomous Systems and the Workshop on Robust Local Descriptors for Computer Vision. LNCS 9009 contains the papers selected for the Workshop on Emerging Topics on Image Restoration and Enhancement, the First International Workshop on Robust Reading, the Second Workshop on User-Centred Computer Vision, the International Workshop on Video Segmentation in Computer Vision, the Workshop: My Car Has Eyes: Intelligent Vehicle with Vision Technology, the Third Workshop on E-Heritage and the Workshop on Computer Vision for Affective Computing. LNCS 9010 contains the papers selected for the Workshop on Feature and Similarity for Computer Vision, the Third International Workshop on Intelligent Mobile and Egocentric Vision and the Workshop on Human Identification for Surveillance.

[Copyright: 246277a877322b384cbbb2cdae5c180b](https://doi.org/10.1007/978-3-319-11967-1)