3d Deep Shape Descriptor Cv Foundation

PIERRE TORRES

Image Analysis Springer
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.
Selected Water Resources Abstracts Springer Nature
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

Scale-dependent/invariant Local 3D Geometric Features and Shape Descriptors John Wiley & Sons
This book presents a detailed review of the state of the art in deep learning approaches for semantic object detection and segmentation in medical image computing, and large-scale radiology database mining. A particular focus is placed on the application of convolutional neural networks, with the theory supported by practical examples. Features: highlights how the use of deep neural networks can address new questions and protocols, as well as improve upon existing challenges in medical image computing; discusses the insightful research experience of Dr. Ronald M. Summers; presents a comprehensive review of the latest research and literature; describes a range of different methods that make use of deep learning for object or landmark detection tasks in 2D and 3D medical imaging; examines a varied selection of techniques for semantic segmentation using deep learning principles in medical imaging; introduces a novel approach to interleaved text and image deep mining on a large-scale radiology image database.
Deep Shape Representations for 3D Object Recognition Springer Nature
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.
Spectral Geometric Methods for Deformable 3D Shape Retrieval Morgan & Claypool Publishers
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
2D and 3D Shape Descriptors Springer Science & Business Media
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.
Molecular Descriptors for Chemoinformatics John Wiley & Sons
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, dep 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.
Intelligence Science and Big Data Engineering. Visual Data Engineering No Starch Press
With the improvement and proliferation of 3D sensors, price cut and enhancementof computational power, the usage of 3D data intensifies for the last few years. The3D point cloud is one type amongst the others for 3D representation. This particularlyrepresentation is the direct output of sensors, accurate and simple. As a non-regularstructure of unordered list of points, the analysis on point cloud is challenging andhence the recent usage only.This PhD thesis focuses on the use of 3D point cloud representation for threedimensional shape analysis. More particularly, the geometrical shape is studied throughthe curvature of the object. Descriptors describing the distribution of the principalcurvature is proposed: Principal Curvature Point Cloud and Multi-Scale PrincipalCurvature Point Cloud. Global Local Point Cloud is another descriptor using thecurvature but in combination with other features. These three descriptors are robustto typical 3D scan error like noisy data or occlusion. They outperform state-of-the-artalgorithms in instance retrieval task with more than 90% of accuracy.The thesis also studies deep learning on 3D point cloud which emerges during thethree years of this PhD. The first approach tested, used curvature-based descriptor asthe input of a multi-layer perceptron network. The accuracy cannot catch state-ofthe-art performances. However, they show that ModelNet, the standard dataset for 3Dshape classification is not a good picture of the reality. Indeed, the experiment showsthat the dataset does not reflect the curvature wealth of true objects scans.Ultimately, a new neural network architecture is proposed. Inspired by the state-ofthe-art deep learning network, Multiscale PointNet computes the feature on multiplescales and combines them all to describe an object. Still under development, theperformances are still to be improved.In summary, tackling the challenging use of 3D point clouds but also the quickevolution of the field, the thesis contributes to the state-of-the-art in three majoraspects: (i) Design of new algorithms, relying on geometrical curvature of the objectfor instance retrieval task. (ii) Study and exhibition of the need to build a new standardclassification dataset with more realistic objects. (iii) Proposition of a new deep neuralnetwork for 3D point cloud analysis.
Computer Vision – ECCV 2020 Springer Nature
□□□□□□□□□□,□□□□□□□□□□□□□□“□”□□□□□,□□□□□□□□□□□□□□□□□□□□.
Programming Computer Vision with Python Springer Nature
This book constitutes the proceedings of the Third International Workshop on Similarity Based Pattern Analysis and Recognition, SIMBAD 2015, which was held in Copenahgen, 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 supervisedto unsupervised learning, from generative to discriminative models, and fromtheoretical issues to empirical validations.
Learning OpenCV 4 Computer Vision with Python 3 Springer Nature
A richly-illustrated, full-color introduction to deep learning that offers visual and conceptual explanations instead of equations. You'll learn how to use key deep learning algorithms without the need for complex math. Ever since computers began beating us at chess, they've been getting better at a wide range of human activities, from writing songs and generating news articles to helping doctors provide healthcare. Deep learning is the source of many of these breakthroughs, and its remarkable ability to find patterns hiding in data has made it the fastest growing field in artificial intelligence (AI). Digital assistants on our phones use deep learning to understand and respond intelligently to voice commands; automotive systems use it to safely navigate road hazards; online platforms use it to deliver personalized suggestions for movies and books - the possibilities are endless. Deep Learning: A Visual Approach is for anyone who wants to understand this fascinating field in depth, but without any of the advanced math and programming usually required to grasp its internals. If you want to know how these tools work, and use them yourself, the answers are all within these pages. And, if you're ready to write your own programs, there are also plenty of supplemental Python notebooks in the accompanying Github repository to get you going. The book's conversational style, extensive color illustrations, illuminating analogies, and real-world examples expertly explain the key concepts in deep learning, including: • How text generators create novel stories and articles • How deep learning systems learn to play and win at human games • How image classification systems identify objects or people in a photo • How to think about probabilities in a way that's useful to everyday life • How to use the machine learning techniques that form the core of modern AI Intellectual adventurers of all kinds can use the powerful ideas covered in Deep Learning: A Visual Approach to build intelligent systems that help us better understand the world and everyone who lives in it. It's the future of AI, and this book allows you to fully envision it. Full Color Illustrations
Applied Mathematics and Computational Mechanics for Smart Applications CRC Press
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.
Sphere Intersection 3D Shape Descriptor (SID) Springer
The availability of large 3D shape benchmarks has sparked a flurry of research activity in the development of efficient techniques for 3D shape recognition, which is a fundamental problem in a variety of domains such as pattern recognition, computer vision, and geometry processing. A key element in virtually any shape recognition method is to represent a 3D shape by a concise and compact shape descriptor aimed at facilitating the recognition tasks. The recent trend in shape recognition is geared toward using deep neural networks to learn features at various levels of abstraction, and has been driven, in large part, by a combination of affordable computing hardware, open source software, and the availability of large-scale datasets. In this thesis, we propose deep learning approaches to 3D shape classification and retrieval. Our approaches inherit many useful properties from the geodesic distance, most notably the capture of the intrinsic geometric structure of 3D shapes and the invariance to isometric deformations. More specifically, we present an integrated framework for 3D shape classification that extracts discriminative geometric shape descriptors with geodesic moments. Further, we introduce a geometric framework for unsupervised 3D shape retrieval using geodesic moments and stacked sparse autoencoders. The key idea is to learn deep shape representations in an unsupervised manner. Such discriminative shape descriptors can then be used to compute pairwise dissimilarities between shapes in a dataset, and to find the retrieved set of the most relevant shapes to a given shape query. Experimental evaluation on three standard 3D shape benchmarks demonstrate the competitive performance of our approach in comparison with existing techniques. We also introduce a deep similarity network fusion framework for 3D shape classification using a graph convolutional neural network, which is an efficient and scalable deep learning model for graph-structured data. The proposed approach coalesces the geometrical discriminative power of geodesic moments and similarity network fusion in an effort to design a simple, yet discriminative shape descriptor. This geometric shape descriptor is then fed into the graph convolutional neural network to learn a deep feature representation of a 3D shape. We validate our method on ModelNet shape benchmarks, demonstrating that the proposed framework yields significant performance gains compared to state-of-the-art approaches.
Descriptor Based Analysis of Digital 3D Shapes "O'Reilly Media, Inc."
Feature descriptors have become a ubiquitous tool in shape analysis. Features can be extracted and subsequently used to design discriminative

signatures for solving a variety of 3D shape analysis problems. In particular, shape classification and retrieval are intriguing and challenging problems that lie at the crossroads of computer vision, geometry processing, machine learning and medical imaging.In this thesis, we propose spectral graph wavelet approaches for the classification and retrieval of deformable 3D shapes. First, we review the recent shape descriptors based on the spectral decomposition of the Laplace-Beltrami operator, which provides a rich set of eigenbases that are invariant to intrinsic isometries. We then provide a detailed overview of spectral graph wavelets. In an effort to capture both local and global characteristics of a 3D shape, we propose a three-step feature description framework. Local descriptors are first extracted via the spectral graph wavelet transform having the Mexican hat wavelet as a generating kernel. Then, mid-level features are obtained by embedding local descriptors into the visual vocabulary space using the soft-assignment coding step of the bag-of-features model. A global descriptor is subsequently constructed by aggregating mid-level features weighted by a geodesic exponential kernel, resulting in a matrix representation that describes the frequency of appearance of nearby codewords in the vocabulary. In order to analyze the performance of the proposed algorithms on 3D shape classification, support vector machines and deep belief networks are applied to mid-level features. To assess the performance of the proposed approach for nonrigid 3D shape retrieval, we compare the global descriptor of a query to the global descriptors of the rest of shapes in the dataset using a dissimilarity measure and find the closest shape. Experimental results on three standard 3D shape benchmarks demonstrate the effectiveness of the proposed classification and retrieval approaches in comparison with state-of-the-art methods.

3D Shape Descriptor Based on 3D Fourier Transform CRC Press

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

Geometric Approaches for 3D Shape Denoising and Retrieval Packt Publishing Ltd

One of the major goals of computer vision is the development of flexible and efficient methods for shape representation. This is true, especially for non-rigid 3D shapes where a great variety of shapes are produced as a result of deformations of a non-rigid object. Modeling these non-rigid shapes is a very challenging problem. Being able to analyze the properties of such shapes and describe their behavior is the key issue in research. Also, considering photometric features can play an important role in many shape analysis applications, such as shape matching and correspondence because it contains rich information about the visual appearance of real objects. This new information (contained in photometric features) and its important applications add another, new dimension to the problem's difficulty. Two main approaches have been adopted in the literature for shape modeling for the matching and retrieval problem, local and global approaches. Local matching is performed between sparse points or regions of the shape, while the global shape approaches similarity is measured among entire models. These methods have an underlying assumption that shapes are rigidly transformed. And Most descriptors proposed so far are confined to shape, that is, they analyze only geometric and/or topological properties of 3D models. A shape descriptor or model should be isometry invariant, scale invariant, be able to capture the fine details of the shape, computationally efficient, and have many other good properties. A shape descriptor or model is needed. This shape descriptor should be: able to deal with the non-rigid shape deformation, able to handle the scale variation problem with less sensitivity to noise, able to match shapes related to the same class even if these shapes have missing parts, and able to encode both the photometric, and geometric information in one descriptor. This dissertation will address the problem of 3D non-rigid shape representation and textured 3D non-rigid shapes based on local features. Two approaches will be proposed for non-rigid shape matching and retrieval based on Heat Kernel (HK), and Scale-Invariant Heat Kernel (SI-HK) and one approach for modeling textured 3D non-rigid shapes based on scale-invariant Weighted Heat Kernel Signature (WHKS). For the first approach, the Laplace-Beltrami eigenfunctions is used to detect a small number of critical points on the shape surface. Then a shape descriptor is formed based on the heat kernels at the detected critical points for different scales. Sparse representation is used to reduce the dimensionality of the calculated descriptor. The proposed descriptor is used for classification via the Collaborative Representation-based Classification with a Regularized Least Square (CRC-RLS) algorithm. The experimental results have shown that the proposed descriptor can achieve state-of-the-art results on two benchmark data sets. For the second approach, an improved method to introduce scale-invariance has been also proposed to avoid noise-sensitive operations in the original transformation method. Then a new 3D shape descriptor is formed based on the histograms of the scale-invariant HK for a number of critical points on the shape at different time scales. A Collaborative Classification (CC) scheme is then employed for object classification. The experimental results have shown that the proposed descriptor can achieve high performance on the two benchmark data sets. An important observation from the experiments is that the proposed approach is more able to handle data under several distortion scenarios (noise, shot-noise, scale, and under missing parts) than the well-known approaches. For modeling textured 3D non-rigid shapes, this dissertation introduces, for the first time, a mathematical framework for the diffusion geometry on textured shapes. This dissertation presents an approach for shape matching and retrieval based on a weighted heat kernel signature. It shows how to include photometric information as a weight over the shape manifold, and it also propose a novel formulation for heat diffusion over weighted manifolds. Then this dissertation presents a new discretization method for the weighted heat kernel induced by the linear FEM weights. Finally, the weighted heat kernel signature is used as a shape descriptor. The proposed descriptor encodes both the photometric, and geometric information based on the solution of one equation. Finally, this dissertation proposes an approach for 3D face recognition based on the front contours of heat propagation over the face surface. The front contours are extracted automatically as heat is propagating starting from a detected set of landmarks. The propagation contours are used to successfully discriminate the various faces. The proposed approach is evaluated on the largest publicly available database of 3D facial images and successfully compared to the state-of-the-art approaches in the literature. This work can be extended to the problem of dense correspondence between non-rigid shapes. The proposed approaches with the properties of the Laplace-Beltrami eigenfunction can be utilized for 3D mesh segmentation. Another possible application of the proposed approach is the view point selection for 3D objects by selecting the most informative views that collectively provide the most descriptive presentation of the surface.

Hyperspectral Image Analysis Springer

The 39-volume set, comprising the LNCS books 13661 until 13699, constitutes the refereed proceedings of the 17th European Conference on Computer Vision, ECCV 2022, held in Tel Aviv, Israel, during October 23–27, 2022. The 1645 papers presented in these proceedings were carefully reviewed and selected from a total of 5804 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.

Similarity-Based Pattern Recognition Springer Nature

This book reviews the state of the art in algorithmic approaches addressing the practical challenges that arise with hyperspectral image analysis tasks, with a focus on emerging trends in machine learning and image processing/understanding. It presents advances in deep learning, multiple instance learning, sparse representation based learning, low-dimensional manifold models, anomalous change detection, target recognition, sensor fusion and super-resolution for robust multispectral and hyperspectral image understanding. It presents research from leading international experts who have made foundational contributions in these areas. The book covers a diverse array of applications of multispectral/hyperspectral imagery in the context of these algorithms, including remote sensing, face recognition and biomedicine. This book would be particularly beneficial to graduate students and researchers who are taking advanced courses in (or are working in) the areas of image analysis, machine learning and remote sensing with multi-channel optical imagery. Researchers and professionals in academia and industry working in areas such as electrical engineering, civil and environmental engineering, geosciences and biomedical image processing, who work with multi-channel optical data will find this book useful.

An Improvement of Rotation Invariant 3D Shape Descriptor Based on Functions on Concentric Spheres Morgan & Claypool Publishers

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.

Mastering OpenCV with Practical Computer Vision Projects Springer Nature

This highly anticipated new edition provides a comprehensive account of face recognition research and technology, spanning the full range of topics needed for designing operational face recognition systems. After a thorough introductory chapter, each of the following chapters focus on a specific topic, reviewing background information, up-to-date techniques, and recent results, as well as offering challenges and future directions. Features: fully updated, revised and expanded, covering the entire spectrum of concepts, methods, and algorithms for automated face detection and recognition systems; provides comprehensive coverage of face detection, tracking, alignment, feature extraction, and recognition technologies, and issues in evaluation, systems, security, and applications; contains numerous step-by-step algorithms; describes a broad range of applications; presents contributions from an international selection of experts; integrates numerous supporting graphs, tables, charts, and performance data.