

Yui Man Lui David Bolme Bruce A Draper J Ross

The two-volume set LNCS 10269 and 10270 constitutes the refereed proceedings of the 20th Scandinavian Conference on Image Analysis, SCIA 2017, held in Tromsø, Norway, in June 2017. The 87 revised papers presented were carefully reviewed and selected from 133 submissions. The contributions are structured in topical sections on history of SCIA; motion analysis and 3D vision; pattern detection and recognition; machine learning; image processing and applications; feature extraction and segmentation; remote sensing; medical and biomedical image analysis; faces, gestures and multispectral analysis.

Foundations of Orientation and Mobility, the classic professional reference and textbook has been completely revised and expanded to two volumes by the most knowledgeable experts in the field. The new third edition includes both the latest research in O&M and expanded information on practice and teaching strategies. Volume 1, History and Theory, includes the bases of O&M knowledge, including perception, orientation, low vision, audition, kinesiology, psychosocial issues, and learning theories, as well as chapters on technology, dog guides, orientation aids, and environmental accessibility. A section on the profession of O&M includes its international history; administration, assessment and program planning; and a chapter on research in O&M. No O&M student or professional can afford to be without this essential resource.

The sixteen-volume set comprising the LNCS volumes 11205-11220 constitutes the refereed proceedings of the 15th European Conference on Computer Vision, ECCV 2018, held in Munich, Germany, in September 2018.The 776 revised papers presented were carefully reviewed and selected from 2439 submissions. The papers are organized in topical sections on learning for vision; computational photography; human analysis; human sensing; stereo and reconstruction; optimization; matching and recognition; video attention; and poster sessions.

This book presents the state of the art in online visual tracking, including the motivations, practical algorithms, and experimental evaluations. Visual tracking remains a highly active area of research in Computer Vision and the performance under complex scenarios has substantially improved, driven by the high demand in connection with real-world applications and the recent advances in machine learning. A large variety of new algorithms have been proposed in the literature over the last two decades, with mixed success. Chapters 1 to 6 introduce readers to tracking methods based on online learning algorithms, including sparse representation, dictionary learning, hashing codes, local model, and model fusion. In Chapter 7, visual tracking is formulated as a foreground/background segmentation problem, and tracking methods based on superpixels and end-to-end deep networks are presented. In turn, Chapters 8 and 9 introduce the cutting-edge tracking methods based on correlation filter and deep learning. Chapter 10 summarizes the book and points out potential future research directions for visual tracking. The book is self-contained and suited for all researchers, professionals and postgraduate students working in the fields of computer vision, pattern recognition, and machine learning. It will help these readers grasp the insights provided by cutting-edge research, and benefit from the practical techniques available for designing effective visual tracking algorithms. Further, the source codes or results of most algorithms in the book are provided at an accompanying website.

Handbook of Biometrics for Forensic Science

Autonomous Vehicles in Support of Naval Operations

Poisons of Plant Origin

Advances in Biometrics

Visual Analysis of Behaviour

In recent years, Moore’s law has fostered the steady growth of the field of digital image processing, though the computational complexity remains a problem for most of the digital image processing applications. In parallel, the research domain of optical image processing has matured, potentially bypassing the problems digital approaches were suffering and bringing new applications. The advancement of technology calls for applications and knowledge at the intersection of both areas but there is a clear knowledge gap between the digital signal processing and the optical processing communities. This book covers the fundamental basis of the optical and image processing techniques by integrating contributions from both optical and digital research communities to solve current application bottlenecks, and give rise to new applications and solutions. Besides focusing on joint research, it also aims at disseminating the knowledge existing in both domains. Applications covered include image restoration, medical imaging, surveillance, holography, etc..." a very good book that deserves to be on the bookshelf of a serious student or scientist working in these areas." Source: Optics and Photonics News

This book presents a comprehensive treatment of visual analysis of behaviour from computational-modelling and algorithm-design perspectives. Topics: covers learning-group activity models, unsupervised behaviour profiling, hierarchical behaviour discovery, learning behavioural context, modelling rare behaviours, and “man-in-the-loop” active learning; examines multi-camera behaviour correlation, person re-identification, and “connecting-the-dots” for abnormal behaviour detection; discusses Bayesian information criterion, Bayesian networks, “bag-of-words” representation, canonical correlation analysis, dynamic Bayesian networks, Gaussian mixtures, and Gibbs sampling; investigates hidden conditional random fields, hidden Markov models, human silhouette shapes, latent Dirichlet allocation, local binary patterns, locality preserving projection, and Markov processes; explores probabilistic graphical models, probabilistic topic models, space-time interest points, spectral clustering, and support vector machines.

Forensic Facial Identification discusses the latest scientific and technical advancements in the field and their implications for practice in psychology, criminology, and law. Provides an up-to-date set of best practices for forensic facial identification Reviews current procedures for different facial identification methods and their reliability Covers eyewitness testimony, line-ups, facial composites, anthropological face reconstructions, CCTV images, and computerized automatic face recognition systems Incorporates case studies which put the latest research and technology in the proper legal context

This comprehensive handbook addresses the sophisticated forensic threats and challenges that have arisen in the modern digital age, and reviews the new computing solutions that have been proposed to tackle them. These include identity-related scenarios which cannot be solved with traditional approaches, such as attacks on security systems and the identification of abnormal/dangerous behaviors from remote cameras. Features: provides an in-depth analysis of the state of the art, together with a broad review of the available technologies and their potential applications; discusses potential future developments in the adoption of advanced technologies for the automated or semi-automated analysis of forensic traces; presents a particular focus on the acquisition and processing of data from real-world forensic cases; offers an holistic perspective, integrating work from different research institutions and combining viewpoints from both biometric technologies and forensic science.

Touchless Fingerprint Biometrics

Perceptual Expertise

15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part V

Video-to-Video Face Recognition for Low-Quality Surveillance Data

Foundations of Indirect Discrimination Law

Optical and Digital Image Processing

The most comprehensive and up-to-date guide to the technologies, applications and human factors considerations of Augmented Reality (AR) and Virtual Reality (VR) systems and wearable computing devices. Practical Augmented Reality is ideal for practitioners and students concerned with any application, from gaming to medicine. It brings together comprehensive coverage of both theory and practice, emphasizing leading-edge displays, sensors, and DIY tools that are already available commercially or will be soon. Beginning with a Foreword by NASA research scientist Victor Luo, this guide begins by explaining the mechanics of human sight, hearing and touch, showing how these perceptual mechanisms (and their performance ranges) directly dictate the design and use of wearable displays, 3-D audio systems, and tactile/force feedback devices. Steve Aukstakalnis presents revealing case studies of real-world applications from gaming, entertainment, science, engineering, aeronautics and aerospace, defense, medicine, telerobotics, architecture, law enforcement, and geophysics. Readers will find clear, easy-to-understand explanations, photos, and illustrations of devices including the Atheer AiR, HTC Vive, DAQRI Smart Helmet, Oculus (Facebook) CV1, Sony PlayStation VR, Vuzix M300, Google Glass, and many more. Functional diagrams and photographs clearly explain how these devices operate, and link directly to relevant theoretical and practical content. Practical Augmented Reality thoroughly considers the human factors of these systems, including sensory and motor physiology constraints, monocular and binocular depth cues, elements contributing to visually-induced motion sickness and nausea, and vergence–accommodation conflicts. It concludes by assessing both the legal and societal implications of new and emerging AR, VR, and wearable technologies as well as provides a look next generation systems.

In Volume 1 of this series, attention was focused on neuropoisons of animal origin. In the present volume, attention has been shifted to poisons of plant origin. In both cases, we have attempted to identify those poisons for which there is a large measure of clinical or research interest. Our efforts in compiling the series have been aided by three groups of individuals. First, we are grateful to the investigators who contributed chapters. Their labors are the substance of this two-volume work. Second, we are pleased to acknowledge the support of Mr. Seymour Weingarten and Plenum Press in our project. And third, we have been immeasurably aided by our assistants, Mrs. Ruby Hough in New York City and Mrs. Helena Walsh in Canberra. To all these persons, we are indebted. Dr. Lance L. Simpson Dr. David R. Curtis New York City Canberra vii Contents Chapter 1 Reserpine 1 by Theodore A. Slot kin I. Introduction

Visual tracking is one of the fundamental problems in computer vision. Its numerous applications include robotics, autonomous driving, augmented reality and 3D reconstruction. In essence, visual tracking can be described as the problem of estimating the trajectory of a target in a sequence of images. The target can be any image region or object of interest. While humans excel at this task, requiring little effort to perform accurate and robust visual tracking, it has proven difficult to automate. It has therefore remained one of the most active research topics in computer vision. In its most general form, no prior knowledge about the object of interest or environment is given, except for the initial target location. This general form of tracking is known as generic visual tracking. The unconstrained nature of this problem makes it particularly difficult, yet applicable to a wider range of scenarios. As no prior knowledge is given, the tracker must learn an appearance model of the target on-the-fly. Cast as a machine learning problem, it imposes several major challenges which are addressed in this thesis. The main purpose of this thesis is the study and advancement of the, so called, Discriminative Correlation Filter (DCF) framework, as it has shown to be particularly suitable for the tracking application. By utilizing properties of the Fourier transform, a correlation filter is discriminatively learned by efficiently minimizing a least-squares objective. The resulting filter is then applied to a new image in order to estimate the target location. This thesis contributes to the advancement of the DCF methodology in several aspects. The main contribution regards the learning of the appearance model: First, the problem of updating the appearance model with new training samples is covered. Efficient update rules and numerical solvers are investigated for this task. Second, the periodic assumption induced by the circular convolution in DCF is countered by proposing a spatial regularization component. Third, an adaptive model of the training set is proposed to alleviate the impact of corrupted or mislabeled training samples. Fourth, a continuous-space formulation of the DCF is introduced, enabling the fusion of multiresolution features and sub-pixel accurate predictions. Finally, the problems of computational complexity and overfitting are addressed by investigating dimensionality reduction techniques. As a second contribution, different feature representations for tracking are investigated. A particular focus is put on the analysis of color features, which had been largely overlooked in prior tracking research. This thesis also studies the use of deep features in DCF-based tracking. While many vision problems have greatly benefited from the advent of deep learning, it has proven difficult to harvest the power of such representations for tracking. In this thesis it is shown that both shallow and deep layers contribute positively. Furthermore, the problem of fusing their complementary properties is investigated. The final major contribution of this thesis regards the prediction of the target scale. In many applications, it is essential to track the scale, or size, of the target since it is strongly related to the relative distance. A thorough analysis of how to integrate scale estimation into the DCF framework is proposed. A one-dimensional scale filter is proposed, enabling efficient and accurate scale estimation.

Summary Oculus Rift in Action introduces the powerful Oculus Rift headset and teaches you how to integrate its many features into 3D games and other virtual reality experiences. You'll start by understanding the capabilities of the Rift hardware. Then you'll follow interesting and instantly-relevant examples that walk you through programming real applications using the Oculus SDK. Examples are provided for both using the Oculus C API directly and for using Unity, a popular development and 3D graphics engine, with the Oculus Unity integration package. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book Virtual reality has long been the domain of researchers and developers with access to specialized hardware and proprietary tools. With the appearance of the Oculus Rift VR headset, the game has changed. Using standard programming tools and the intuitive Oculus SDKs, you can deliver powerful immersive games, simulations, and other virtual experiences that finally nail the feeling of being in the middle of the action. Oculus Rift in Action teaches you how to create 3D games and other virtual reality experiences for the Oculus Rift. You'll explore the Rift hardware through examples of real applications using the Oculus SDK and both the Oculus C API and the Unity 3D graphics engine. Along the way, you'll get practical guidance on how to use the Rift's sensors to produce fluid VR experiences. Experience with C++, C#, or another OO language is assumed. What's Inside Creating immersive VR Integrating the Rift with the Unity 3D SDK Implementing the mathematics of 3D Avoiding motion-sickness triggers About the Authors Brad Davis is an active VR developer who maintains a great set of example Rift applications on Github. Karen Bryla is a freelance developer and writer. Alex Benton is a lecturer in 3D graphics at the University of Cambridge and a software engineer at Google. Table of Contents PART 1 GETTING STARTED Meet the Oculus Rift PART 2 USING THE OCULUS C API Creating your first Rift interactions Pulling data out of the Rift: working with the head tracker Sending output to the Rift: working with the display Putting it all together: integrating head tracking and 3D rendering Performance and quality PART 3 USING UNITY: creating applications that run on the Rift Unity: tailoring your application for the Rift PART 4 THE VR USER EXPERIENCE UI design for VR Reducing motion sickness and discomfort PART 5 ADVANCED RIFT INTEGRATIONS Using the Rift with Java and Python Case study: a VR shader editor Augmenting virtual reality

The True Story of a Girl Caught Between Two Worlds

7th International Conference on Computer Vision Systems, ICVS 2009 Liège, Belgium, October 13-15, 2009, Proceedings

Learning Convolution Operators for Visual Tracking

The Last Feather

Handbook of Iris Recognition

Statistical Methods for Performance Evaluation

Biometrics, the science of using physical traits to identify individuals, is playing an increasing role in our security-conscious society and across the globe. Biometric authentication, or bioauthentication, systems are being used to secure everything from amusement parks to bank accounts to military installations. Yet developments in this field have not been matched by an equivalent improvement in the statistical methods for evaluating these systems. Compensating for this need, this unique text/reference provides a basic statistical methodology for practitioners and testers of bioauthentication devices, supplying a set of rigorous statistical methods for evaluating biometric authentication systems. This framework of methods can be extended and generalized for a wide range of applications and tests. This is the first single resource on statistical methods for estimation and comparison of the performance of biometric authentication systems. The book focuses on six common performance metrics: for each metric, statistical methods are derived for a single system that incorporates confidence intervals, hypothesis tests, sample size calculations, power calculations and prediction intervals. These methods are also extended to allow for the statistical comparison and evaluation of multiple systems for both independent and paired data. Topics and features: * Provides a statistical methodology for the most common biometric performance metrics: failure to enroll (FTE), failure to acquire (FTA), false non-match rate (FNMR), false match rate (FMR), and receiver operating characteristic (ROC) curves * Presents methods for the comparison of two or more biometric performance metrics * Introduces a new bootstrap methodology for FMR and ROC curve estimation * Supplies more than 120 examples, using publicly available biometric data where possible * Discusses the addition of prediction intervals to the bioauthentication statistical toolset * Describes sample-size and power calculations for FTE, FTA, FNMR and FMR Researchers, managers and decisions makers needing to compare biometric systems across a variety of metrics will find within this reference an invaluable set of statistical tools. Written for an upper-level undergraduate or master’s level audience with a quantitative background, readers are also expected to have an understanding of the topics in a typical undergraduate statistics course. Dr. Michael E. Schuckers is Associate Professor of Statistics at St. Lawrence University, Canton, NY, and a member of the Center for Identification Technology Research.

This book reviews cutting-edge developments in neural signalling processing (NSP), systematically introducing readers to various models and methods in the context of NSP. Neuronal Signal Processing is a comparatively new field in computer sciences and neuroscience, and is rapidly establishing itself as an important tool, one that offers an ideal opportunity to forge stronger links between experimentalists and computer scientists. This new signal-processing tool can be used in conjunction with existing computational tools to analyse neural activity, which is monitored through different sensors such as spike trains, local filed potentials and EEG. The analysis of neural activity can yield vital insights into the function of the brain. This book highlights the contribution of signal processing in the area of computational neuroscience by providing a forum for researchers in this field to share their experiences to date.

Video Tracking provides a comprehensive treatment of the fundamental aspects of algorithm and application development for the task of estimating, over time, the position of objects of interest seen through cameras. Starting from the general problem definition and a review of existing and emerging video tracking applications, the book discusses popular methods, such as those based on correlation and gradient-descent. Using practical examples, the reader is introduced to the advantages and limitations of deterministic approaches, and is then guided toward more advanced video tracking solutions, such as those based on the Bayes’ recursive framework and on Random Finite Sets. Key features: Discusses the design choices and implementation issues required to turn the underlying mathematical models into a real-world effective tracking systems. Provides block diagrams and simil-code implementation of the algorithms. Reviews methods to evaluate the performance of video trackers – this is identified as a major problem by end-users. The book aims to help researchers and practitioners develop techniques and solutions based on the potential of video tracking applications. The design methodologies discussed throughout the book provide guidelines for developers in the industry working on vision-based applications. The book may also serve as a reference for engineering and computer science graduate students involved in vision, robotics, human-computer interaction, smart environments and virtual reality programmes

Offering the first comprehensive analysis of touchless fingerprint-recognition technologies, Touchless Fingerprint Biometrics gives an overview of the state of the art and describes relevant industrial applications. It also presents new techniques to efficiently and effectively implement advanced solutions based on touchless fingerprinting.The most

Handbook of Biometric Anti-Spoofing

From Pixels to Semantics

An Analysis of Current and Future Technology for Unmanned Maritime Vehicles

High Energy Density Materials

Theory and Practice

Signal Processing in Neuroscience

A practical guide designed to get you from basics to current state of art in computer vision systems. Key Features Master the different tasks associated with Computer Vision and develop your own Computer Vision applications with ease Leverage the power of Python, Tensorflow, Keras, and OpenCV to perform image processing, object detection, feature detection and more With real-world datasets and fully functional code, this book is your one-stop guide to understanding Computer Vision Book Description In this book, you will find several recently proposed methods in various domains of computer vision. You will start by setting up the proper Python environment to work on practical applications. This includes setting up libraries such as OpenCV, TensorFlow, and Keras using Anaconda. Using these libraries, you'll start to understand the concepts of image transformation and filtering. You will find a detailed explanation of feature detectors such as FAST and ORB; you'll use them to find similar-looking objects. With an introduction to convolutional neural nets, you will learn how to build a deep neural net using Keras and how to use it to classify the Fashion-MNIST dataset. With regard to object detection, you will learn the implementation of a simple face detector as well as the workings of complex deep-learning-based object detectors such as Faster R-CNN and SSD using TensorFlow. You'll get started with semantic segmentation using FCN models and track objects with Deep SORT. Not only this, you will also use Visual SLAM techniques such as ORB-SLAM on a standard dataset. By the end of this book, you will have a firm understanding of the different computer vision techniques and how to apply them in your applications. What you will learn Learn the basics of image manipulation with OpenCV Implement and visualize image filters such as smoothing, dilation, histogram equalization, and more Set up various libraries and platforms, such as OpenCV, Keras, and TensorFlow, in order to start using computer vision, along with appropriate datasets for each chapter, such as MSCOCO, MOT, and Fashion-MNIST Understand image transformation and downsampling with practical implementations. Explore neural networks for computer vision and convolutional neural networks using Keras Understand working on deep-learning-based object detection such as Faster-R-CNN, SSD, and more Explore deep-learning-based object tracking in action Understand Visual SLAM techniques such as ORB-SLAM Who this book is for This book is for machine learning practitioners and deep learning enthusiasts who want to understand and implement various tasks associated with Computer Vision and image processing in the most practical manner possible. Some programming experience would be beneficial while knowing Python would be an added bonus.

