Video Analysis Deep Learning

Granular Video Computing
Learning OpenCV 3 Computer Vision with Python
Video and Image Analysis Using Statistical and Machine Learning Techniques
Image Analysis and Recognition
Deep Learning for Action Understanding in Video
Video Analytics Using Deep Learning
3D Imaging—Multidimensional Signal Processing and Deep Learning
Intelligent Image and Video Analytics
Deep Learning for Medical Applications with Unique Data
Deep Learning for Computer Vision
Mastering Computer Vision with TensorFlow 2.x
Topic-based Video Classification and Retrieval Using Machine Learning
Granular Video Computing: With Rough Sets, Deep Learning And In Iot
Deep Learning Applications in Image Analysis
Machine Learning for Human Motion Analysis: Theory and Practice
Efficient Processing of Deep Neural Networks
Machine Learning for Big Data Analysis
Image Analysis and Processing – ICIAP 2019
Image Analysis and Processing – ICIAP 2022
Deep Learning for Robot Perception and Cognition
Explainable Deep Learning AI
Machine Learning for Audio, Image and Video Analysis
Overcomplete Dictionary and Deep Learning Approaches to Image and Video Analysis
Semi-supervised Deep Learning with Applications in Surgical Video Analysis and Bioinformatics
Deep-Learning-based Video Analysis for Human Action Evaluation
Deep Learning for Coders with fastai and PyTorch
Roadside Video Data Analysis
Image Analysis and Processing – ICIAP 2019
Structured Deep Learning for Video Analysis
Codeless Deep Learning with KNIME
Advanced Image and Video Processing Using MATLAB
Machine Learning Methods for Behaviour Analysis and Anomaly Detection in Video
Bridging the Semantic Gap in Image and Video Analysis
Learning Hierarchical Representations for Video Analysis Using Deep Learning
Visual Object Tracking from Correlation Filter to Deep Learning
Proceedings of 3rd International Conference on Computer Vision and Image Processing
Data Science and Deep Learning Workshop For Scientists and Engineers
Deep Learning in Mining of Visual Content

Video Analysis Deep Learning

BROOKS MYLA

Granular Video Computing CRC Press
"This book highlights the development of robust and effective vision-based motion understanding systems, addressing specific vision applications such as surveillance, sport event analysis, healthcare, video conferencing, and motion video indexing and retrieval"--Provided by publisher.
Learning OpenCV 3 Computer Vision with Python Walter de Gruyter GmbH & Co KG
With the massive increase of video content on Internet and beyond, the automatic understanding of visual content could impact many different application fields such as robotics, health care, content search or filtering. The goal of this thesis is to provide methodological contributions in Computer Vision and Machine Learning for automatic content understanding from videos. We emphasis on problems, namely fine-grained human action recognition and visual reasoning from object-level interactions. In the first part of this manuscript, we tackle the problem of fine-grained human action recognition. We introduce two different trained attention mechanisms on the visual content from articulated human pose. The first method is able to automatically draw attention to important pre-selected points of the video conditioned on learned features extracted from the articulated human pose. We show that such mechanism improves performance on the final task and provides a good way to visualize the most discriminative parts of the visual content. The second method goes beyond pose-based human action recognition. We develop a method able to automatically identify unstructured feature clouds of interest in the video using contextual information. Furthermore, we introduce a learned distributed system for aggregating the features in a recurrent manner and taking decisions in a distributed way. We demonstrate that we can achieve a better performance than obtained previously, without using articulated pose information at test time. In the second part of this thesis, we investigate video representations from an object-level perspective. Given a set of detected persons and objects in the scene, we develop a method which learns to infer the important object interactions through space and time using the video-level annotation only. That allows to identify important objects and object interactions for a given action, as well as potential dataset bias. Finally, in a third part, we go beyond the task of classification and supervised learning from visual content by tackling causality in interactions, in particular the problem of counterfactual learning. We introduce a new benchmark, namely CoPhy, where, after watching a video, the task is to predict the outcome after modifying the initial stage of the video. We develop a method based on object- level interactions able to infer object properties without supervision as well as future object locations after the intervention.
Video and Image Analysis Using Statistical and Machine Learning Techniques Springer
Explainable Deep Learning AI: Methods and Challenges presents the latest works of leading researchers in the XAI area, offering an overview of the XAI area, along with several novel technical methods and applications that address explainability challenges for deep learning AI systems. The book overviews XAI and then covers a number of specific technical works and approaches for deep learning, ranging from general XAI methods to specific XAI applications, and finally, with user-oriented evaluation approaches. It also explores the main categories of explainable AI – deep learning, which become the necessary condition in various applications of artificial intelligence. The groups of methods such as back-propagation and perturbation-based methods are explained, and the application to various kinds of data classification are presented. Provides an overview of main approaches to Explainable Artificial Intelligence (XAI) in the Deep Learning realm, including the most popular techniques and their use, concluding with challenges and exciting future directions of XAI Explores the latest developments in general XAI methods for Deep Learning Explains how XAI for Deep Learning is applied to various domains like images, medicine and natural language processing Provides an overview of how XAI systems are tested and evaluated, specially with real users, a critical need in XAI
Image Analysis and Recognition Springer Nature
Discover how to integrate KNIME Analytics Platform with deep learning libraries to implement artificial intelligence solutions Key FeaturesBecome well-versed with KNIME Analytics Platform to perform codeless deep learningDesign and build deep learning workflows quickly and more easily using the KNIME GUIDiscover different deployment options without using a single line of code with KNIME Analytics PlatformBook Description KNIME Analytics Platform is an open source software used to create and design data science workflows. This book is a comprehensive guide to the KNIME GUI and KNIME deep learning integration, helping you build neural network models without writing any code. It'll guide you in building simple and complex neural networks through practical and creative solutions for solving real-world data problems. Starting with an introduction to KNIME Analytics Platform, you'll get an overview of simple feed-forward networks for solving simple classification problems on relatively small datasets. You'll then move on to build, train, test, and deploy more complex networks, such as autoencoders, recurrent neural networks (RNNs), long short-term memory (LSTM), and convolutional neural networks (CNNs). In each chapter, depending on the network and use case, you'll learn how to prepare data, encode incoming data, and apply best practices. By the end of this book, you'll have learned how to design a variety of different neural architectures and will be able to train, test, and deploy the final network. What you will learnUse various common nodes to transform your data into the right structure suitable for training a neural networkUnderstand neural network techniques such as loss functions, backpropagation, and hyperparametersPrepare and encode data appropriately to feed it into the networkBuild and train a classic feedforward networkDevelop and optimize an autoencoder network for outlier detectionImplement deep learning networks such as CNNs, RNNs, and LSTM with the help of practical examplesDeploy a trained deep learning network on real-world dataWho this book is for This book is for data analysts, data scientists, and deep learning developers who are not well-versed in Python but want to learn how to use KNIME GUI to build, train, test, and deploy neural networks with different architectures. The practical implementations shown in the book do not require coding or any knowledge of dedicated scripts, so you can easily implement your knowledge into practical

applications. No prior experience of using KNIME is required to get started with this book.

Deep Learning for Action Understanding in Video Springer Nature

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve key metrics-such as energy-efficiency, throughput, and latency-without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as formalization and organization of key concepts from contemporary work that provide insights that may spark new ideas.

Video Analytics Using Deep Learning Springer

This thesis proposes machine learning methods for understanding scenes via behaviour analysis and online anomaly detection in video. The book introduces novel Bayesian topic models for detection of events that are different from typical activities and a novel framework for change point detection for identifying sudden behavioural changes. Behaviour analysis and anomaly detection are key components of intelligent vision systems. Anomaly detection can be considered from two perspectives: abnormal events can be defined as those that violate typical activities or as a sudden change in behaviour. Topic modelling and change-point detection methodologies, respectively, are employed to achieve these objectives. The thesis starts with the development of learning algorithms for a dynamic topic model, which extract topics that represent typical activities of a scene. These typical activities are used in a normality measure in anomaly detection decision-making. The book also proposes a novel anomaly localisation procedure. In the first topic model presented, a number of topics should be specified in advance. A novel dynamic nonparametric hierarchical Dirichlet process topic model is then developed where the number of topics is determined from data. Batch and online inference algorithms are developed. The latter part of the thesis considers behaviour analysis and anomaly detection within the change-point detection methodology. A novel general framework for change-point detection is introduced. Gaussian process time series data is considered. Statistical hypothesis tests are proposed for both offline and online data processing and multiple change point detection are proposed and theoretical properties of the tests are derived. The thesis is accompanied by open-source toolboxes that can be used by researchers and engineers.

3D Imaging—Multidimensional Signal Processing and Deep Learning Springer

This volume links the concept of granular computing using deep learning and the Internet of Things to object tracking for video analysis. It describes how uncertainties, involved in the task of video processing, could be handled in rough set theoretic granular computing frameworks. Issues such as object tracking from videos in constrained situations, occlusion/overlapping handling, measuring of the reliability of tracking methods, object recognition and linguistic interpretation in video scenes, and event prediction from videos, are the addressed in this volume. The book also looks at ways to reduce data dependency in the context of unsupervised (without manual interaction/ labeled data/ prior information) training.This book may be used both as a textbook and reference book for graduate students and researchers in computer science, electrical engineering, system science, data science, and information technology, and is recommended for both students and practitioners working in computer vision, machine learning, video analytics, image analytics, artificial intelligence, system design, rough set theory, granular computing, and soft computing.

Intelligent Image and Video Analytics Elsevier

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Deep Learning for Medical Applications with Unique Data "O'Reilly Media, Inc."

This book presents cutting-edge research on various ways to bridge the semantic gap in image and video analysis. The respective chapters address different stages of image processing, revealing that the first step is a future extraction, the second is a segmentation process, the third is object recognition, and the fourth and last involve the semantic interpretation of the image. The semantic gap is a challenging area of research, and describes the difference between low-level features extracted from the image and the high-level semantic meanings that people can derive from the image. The result greatly depends on lower level vision techniques, such as feature selection, segmentation, object recognition, and so on. The use of deep models has freed humans from manually selecting and extracting the set of features. Deep learning does this automatically, developing more abstract features at the successive levels. The book offers a valuable resource for researchers, practitioners, students and professors in Computer Engineering, Computer Science and related fields whose work involves images, video analysis, image interpretation and so on.

Deep Learning for Computer Vision Packt Publishing Ltd

This book offers a comprehensive introduction to advanced methods for image and video analysis and processing. It covers deraining, dehazing, inpainting, fusion, watermarking and stitching. It describes techniques for face and lip recognition,

facial expression recognition, lip reading in videos, moving object tracking, dynamic scene classification, among others. The book combines the latest machine learning methods with computer vision applications, covering topics such as event recognition based on deep learning,dynamic scene classification based on topic model, person re-identification based on metric learning and behavior analysis. It also offers a systematic introduction to image evaluation criteria showing how to use them in different experimental contexts. The book offers an example-based practical guide to researchers, professionals and graduate students dealing with advanced problems in image analysis and computer vision.

Mastering Computer Vision with TensorFlow 2.x Springer Nature
Digital videos and images are effective media for capturing spatial and temporal information in the real world. The rapid growth of digital videos has motivated research aimed at developing effective algorithms, with the objective of obtaining useful information for a variety of application areas, such as security, commerce, medicine, geography, etc. This dissertation presents innovative and practical techniques, based on statistics and machine learning, that address some key research problems in video and image analysis, including video stabilization, object classification, image segmentation, and video indexing.

Topic-based Video Classification and Retrieval Using Machine Learning Springer Nature
On three action recognition benchmarks, namely HMDB-51, UCF101 and a subset of Kinetics, we demonstrate that our DMC-Net can significantly shorten the performance gap between state-of-the-art compressed video based methods with and without optical flow, while being two orders of magnitude faster than the methods that use optical flow. By addressing the three major challenges mentioned above, we are able to develop more robust models for video action understanding and improve performance in various dimensions, such as (1) temporal precision, (2) required levels of supervision, (3) live video analysis ability, and finally (4) efficiency in processing compressed video. Our research has contributed significantly to advancing the state of the art of video action understanding and expanding the foundation for comprehensive semantic understanding of video content.

Granular Video Computing: With Rough Sets, Deep Learning And In Iot Packt Publishing Ltd
Learn how to model and train advanced neural networks to implement a variety of Computer Vision tasks Key Features Train different kinds of deep learning model from scratch to solve specific problems in Computer Vision Combine the power of Python, Keras, and TensorFlow to build deep learning models for object detection, image classification, similarity learning, image captioning, and more Includes tips on optimizing and improving the performance of your models under various constraints Book Description Deep learning has shown its power in several application areas of Artificial Intelligence, especially in Computer Vision. Computer Vision is the science of understanding and manipulating images, and finds enormous applications in the areas of robotics, automation, and so on. This book will also show you, with practical examples, how to develop Computer Vision applications by leveraging the power of deep learning. In this book, you will learn different techniques related to object classification, object detection, image segmentation, captioning, image generation, face analysis, and more. You will also explore their applications using popular Python libraries such as TensorFlow and Keras. This book will help you master state-of-the-art, deep learning algorithms and their implementation. What you will learn Set up an environment for deep learning with Python, TensorFlow, and Keras Define and train a model for image and video classification Use features from a pre-trained Convolutional Neural Network model for image retrieval Understand and implement object detection using the real-world Pedestrian Detection scenario Learn about various problems in image captioning and how to overcome them by training images and text together Implement similarity matching and train a model for face recognition Understand the concept of generative models and use them for image generation Deploy your deep learning models and optimize them for high performance Who this book is for This book is targeted at data scientists and Computer Vision practitioners who wish to apply the concepts of Deep Learning to overcome any problem related to Computer Vision. A basic knowledge of programming in Python—and some understanding of machine learning concepts—is required to get the best out of this book.

Deep Learning Applications in Image Analysis Apress
The book focuses on visual object tracking systems and approaches based on correlation filter and deep learning. Both foundations and implementations have been addressed. The algorithm, system design and performance evaluation have been explored for three kinds of tracking methods including correlation filter based methods, correlation filter with deep feature based methods, and deep learning based methods. Firstly, context aware and multi-scale strategy are presented in correlation filter based trackers; then, long-short term correlation filter, context aware correlation filter and auxiliary relocation in SiamFC framework are proposed for combining correlation filter and deep learning in visual object tracking; finally, improvements in deep learning based trackers including Siamese network, GAN and reinforcement learning are designed. The goal of this book is to bring, in a timely fashion, the latest advances and developments in visual object tracking, especially correlation filter and deep learning based methods, which is particularly suited for readers who are interested in the research and technology innovation in visual object tracking and related fields.

Machine Learning for Human Motion Analysis: Theory and Practice Academic Press
The two-volume set LNCS 11751 and 11752 constitutes the refereed proceedings of the 20th International Conference on Image Analysis and Processing, ICIAP 2019, held in Trento, Italy, in September 2019. The 117 papers presented were carefully reviewed and selected from 207 submissions. The papers cover both classic and the most recent trends in image processing, computer vision, and pattern recognition, addressing both theoretical and applicative aspects. They are organized in the following topical sections: Video Analysis and Understanding; Pattern Recognition and Machine Learning; Deep Learning; Multiview Geometry and 3D Computer Vision; Image Analysis, Detection and Recognition; Multimedia; Biomedical and Assistive Technology; Digital Forensics; Image processing for Cultural Heritage.

Efficient Processing of Deep Neural Networks Packt Publishing Ltd
Deep Learning for Medical Applications with Unique Data informs readers about the most recent deep learning-based medical applications in which only unique data gathered in real cases are used. The book provides examples of how deep learning can be used in different problem areas and frameworks in both clinical and research settings, including medical image analysis, medical image registration, time series analysis, medical data synthesis,

drug discovery, and pre-processing operations. The volume discusses not only positive findings, but also negative ones obtained by deep learning techniques, including the use of newly developed deep learning techniques rarely reported in the existing literature. The book excludes research works with ready data sets and includes only unique data use to better understand the state of deep learning in real-world cases, along with the feedback and user experiences from physicians and medical staff for applied deep learning-based solutions. Other applications presented in the book include hybrid solutions with deep learning support, disease diagnosis with deep learning focusing on rare diseases and cancer, patient care and treatment, genomics research, as well as research on robotics and autonomous systems. Introduces deep learning, demonstrating concepts for a wide variety of medical applications using unique data, excluding research with ready datasets Encompasses a wide variety of biomedical applications, including unsupervised learning, natural language processing, pattern recognition, image and video processing and disease diagnosis Provides a robust set of methods that will help readers appropriately and judiciously use the most suitable deep learning techniques for their applications

Machine Learning for Big Data Analysis Springer Nature
The proceedings set LNCS 13231, 13232, and 13233 constitutes the refereed proceedings of the 21st International Conference on Image Analysis and Processing, ICIAP 2022, which was held during May 23-27, 2022, in Lecce, Italy, The 168 papers included in the proceedings were carefully reviewed and selected from 307 submissions. They deal with video analysis and understanding; pattern recognition and machine learning; deep learning; multi-view geometry and 3D computer vision; image analysis, detection and recognition; multimedia; biomedical and assistive technology; digital forensics and biometrics; image processing for cultural heritage; robot vision; etc.

Image Analysis and Processing – ICIAP 2019 Springer Nature
As video analysis provides an automatic solution to extract meaningful information from the video content, it can be applied in healthcare to evaluate human action patterns for various purposes, such as biometrics estimation and performance assessment. In recent years, the fast development of deep learning and portable medical sensors has led to more affordable and accurate computer vision-based measurements for human action patterns, thus enabling a more efficient video analysis system for action evaluation in home and clinic environments. We investigate the novel usage of video analysis for healthcare monitoring purposes, including objective biometrics estimation and subjective action quality assessment. We propose a deep learning framework to extract spatial-temporal features and estimate biometrics or performance scores from 3D body landmarks using a graph convolutional neural network, which offers a portable solution to obtain gold-standard biometrics with 3D multi-joint coordination underlying body movements and can provide real-time feedback of movement performance for rehabilitation exercises. For biometrics estimation, in Chapter 2, we propose two single-task models for video-level and frame-level estimation, respectively, and a multi-task learning approach to estimate CoP metrics on two different temporal levels in parallel. To facilitate this line of research, we collect and release a novel computer-vision-based 3D body landmark dataset using pose estimation. We extend our framework to a traditional kinematics dataset collected by on-body reflective markers by using adaptive graph convolution. For action quality assessment, we propose a deep learning framework for automatic assessment of physical rehabilitation exercises using a graph convolutional network with self-supervised regularization in Chapter 3. To further improve the accessibility of the real-time CoP metrics estimation system, we investigate a view-invariant video-level CoP metrics estimation framework using a single RGB camera in Chapter 4, which could significantly benefit the data collection in home and clinic environments. We also explore a semi-supervised learning framework for video-level CoP metrics estimation for partially labeled data with only a small portion of labels in Chapter 5. Our proposed methods potentially enable a more affordable, comprehensive, and portable virtual therapy system than is available with existing tools.

Image Analysis and Processing – ICIAP 2022 World Scientific
Deep Learning for Robot Perception and Cognition introduces a broad range of topics and methods in deep learning for robot perception and cognition together with end-to-end methodologies. The book provides the conceptual and mathematical background needed for approaching a large number of robot perception and cognition tasks from an end-to-end learning point-of-view. The book is suitable for students, university and industry researchers and practitioners in Robotic Vision, Intelligent Control, Mechatronics, Deep Learning, Robotic Perception and Cognition tasks. Presents deep learning principles and methodologies Explains the principles of applying end-to-end learning in robotics applications Presents how to design and train deep learning models Shows how to apply deep learning in robot vision tasks such as object recognition, image classification, video analysis, and more Uses robotic simulation environments for training deep learning models Applies deep learning methods for different tasks ranging from planning and navigation to biosignal analysis

Deep Learning for Robot Perception and Cognition
Academic Press
This second edition focuses on audio, image and video data, the three main types of input that machines deal with when interacting with the real world. A set of appendices provides the reader with self-contained introductions to the mathematical background necessary to read the book. Divided into three main parts, From Perception to Computation introduces methodologies aimed at representing the data in forms suitable for computer processing, especially when it comes to audio and images. Whilst the second part, Machine Learning includes an extensive overview of statistical techniques aimed at addressing three main problems, namely classification (automatically assigning a data sample to one of the classes belonging to a predefined set), clustering (automatically grouping data samples according to the similarity of their properties) and sequence analysis (automatically mapping a sequence of observations into a sequence of human-understandable symbols). The third part Applications shows how the abstract problems defined in the second part underlie technologies capable to perform complex tasks such as the recognition of hand gestures or the transcription of handwritten data. Machine Learning for Audio, Image and Video Analysis is suitable for students to acquire a solid background in machine learning as well as for practitioners to deepen their knowledge of the state-of-the-art. All application chapters are based on publicly available data and free software packages, thus allowing readers to replicate the experiments.

Related with Video Analysis Deep Learning:
© Video Analysis Deep Learning Free Anatomy And Physiology Quiz
© Video Analysis Deep Learning Free Curriculum Mapping Templates
© Video Analysis Deep Learning Free Asp Practice Test