Lecture Notes in Computer Science **11132**

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison
Lancaster University, Lancaster, UK
Takeo Kanade
Carnegie Mellon University, Pittsburgh, PA, USA
Josef Kittler
University of Surrey, Guildford, UK
Jon M. Kleinberg
Cornell University, Ithaca, NY, USA
Friedemann Mattern
ETH Zurich, Zurich, Switzerland
John C. Mitchell
Stanford University, Stanford, CA, USA
Moni Naor
Weizmann Institute of Science, Rehovot, Israel
C. Pandu Rangan
Indian Institute of Technology Madras, Chennai, India
Bernhard Steffen
TU Dortmund University, Dortmund, Germany
Demetri Terzopoulos
University of California, Los Angeles, CA, USA
Doug Tygar
University of California, Berkeley, CA, USA

More information about this series at http://www.springer.com/series/7412

Laura Leal-Taixé · Stefan Roth (Eds.)

Computer Vision – ECCV 2018 Workshops

Munich, Germany, September 8–14, 2018
Proceedings, Part IV

Springer

Editors
Laura Leal-Taixé
Technical University of Munich
Garching, Germany

Stefan Roth
Technische Universität Darmstadt
Darmstadt, Germany

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-11017-8 ISBN 978-3-030-11018-5 (eBook)
https://doi.org/10.1007/978-3-030-11018-5

Library of Congress Control Number: 2018966826

LNCS Sublibrary: SL6 – Image Processing, Computer Vision, Pattern Recognition, and Graphics

© Springer Nature Switzerland AG 2019
This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword

It was our great pleasure to host the European Conference on Computer Vision 2018 in Munich, Germany. This constituted by far the largest ECCV event ever. With close to 2,900 registered participants and another 600 on the waiting list one month before the conference, participation more than doubled since the last ECCV in Amsterdam. We believe that this is due to a dramatic growth of the computer vision community combined with the popularity of Munich as a major European hub of culture, science, and industry. The conference took place in the heart of Munich in the concert hall Gasteig with workshops and tutorials held on the downtown campus of the Technical University of Munich.

One of the major innovations for ECCV 2018 was the free perpetual availability of all conference and workshop papers, which is often referred to as open access. We note that this is not precisely the same use of the term as in the Budapest declaration. Since 2013, CVPR and ICCV have had their papers hosted by the Computer Vision Foundation (CVF), in parallel with the IEEE Xplore version. This has proved highly beneficial to the computer vision community.

We are delighted to announce that for ECCV 2018 a very similar arrangement was put in place with the cooperation of Springer. In particular, the author's final version will be freely available in perpetuity on a CVF page, while SpringerLink will continue to host a version with further improvements, such as activating reference links and including video. We believe that this will give readers the best of both worlds; researchers who are focused on the technical content will have a freely available version in an easily accessible place, while subscribers to SpringerLink will continue to have the additional benefits that this provides. We thank Alfred Hofmann from Springer for helping to negotiate this agreement, which we expect will continue for future versions of ECCV.

September 2018

Horst Bischof
Daniel Cremers
Bernt Schiele
Ramin Zabih

Preface

It is our great pleasure to present these workshop proceedings of the 15th European Conference on Computer Vision, which was held during September 8–14, 2018, in Munich, Germany. We are delighted that the main conference of ECCV 2018 was accompanied by 43 scientific workshops. The ECCV workshop proceedings contain contributions of 36 workshops.

We received 74 workshop proposals on a broad set of topics related to computer vision. The very high quality and the large number of proposals made the selection process rather challenging. Owing to space restrictions, only 46 proposals were accepted, among which six proposals were merged into three workshops because of overlapping themes.

The final set of 43 workshops complemented the main conference program well. The workshop topics presented a good orchestration of new trends and traditional issues, built bridges into neighboring fields, as well as discussed fundamental technologies and novel applications. We would like to thank all the workshop organizers for their unreserved efforts to make the workshop sessions a great success.

September 2018

Stefan Roth
Laura Leal-Taixé

Organization

General Chairs

Horst Bischof	Graz University of Technology, Austria
Daniel Cremers	Technical University of Munich, Germany
Bernt Schiele	Saarland University, Max Planck Institute for Informatics, Germany
Ramin Zabih	CornellNYCTech, USA

Program Chairs

Vittorio Ferrari	University of Edinburgh, UK
Martial Hebert	Carnegie Mellon University, USA
Cristian Sminchisescu	Lund University, Sweden
Yair Weiss	Hebrew University, Israel

Local Arrangement Chairs

Björn Menze	Technical University of Munich, Germany
Matthias Niessner	Technical University of Munich, Germany

Workshop Chairs

Stefan Roth	Technische Universität Darmstadt, Germany
Laura Leal-Taixé	Technical University of Munich, Germany

Tutorial Chairs

Michael Bronstein	Università della Svizzera Italiana, Switzerland
Laura Leal-Taixé	Technical University of Munich, Germany

Website Chair

Friedrich Fraundorfer	Graz University of Technology, Austria

Demo Chairs

Federico Tombari	Technical University of Munich, Germany
Joerg Stueckler	Technical University of Munich, Germany

Publicity Chair

Giovanni Maria Farinella	University of Catania, Italy

Industrial Liaison Chairs

Florent Perronnin	Naver Labs, France
Yunchao Gong	Snap, USA
Helmut Grabner	Logitech, Switzerland

Finance Chair

Gerard Medioni	Amazon, University of Southern California, USA

Publication Chairs

Albert Ali Salah	Boğaziçi University, Turkey
Hamdi Dibeklioğlu	Bilkent University, Turkey
Anton Milan	Amazon, Germany

Workshop Organizers

W01 – The Visual Object Tracking Challenge Workshop

Matej Kristan	University of Ljubljana, Slovenia
Aleš Leonardis	University of Birmingham, UK
Jiří Matas	Czech Technical University in Prague, Czechia
Michael Felsberg	Linköping University, Sweden
Roman Pflugfelder	Austrian Institute of Technology, Austria

W02 – 6th Workshop on Computer Vision for Road Scene Understanding and Autonomous Driving

Mathieu Salzmann	EPFL, Switzerland
José Alvarez	NVIDIA, USA
Lars Petersson	Data61 CSIRO, Australia
Fredrik Kahl	Chalmers University of Technology, Sweden
Bart Nabbe	Aurora, USA

W03 – 3D Reconstruction in the Wild

Akihiro Sugimoto	The National Institute of Informatics (NII), Japan
Tomas Pajdla	Czech Technical University in Prague, Czechia
Takeshi Masuda	The National Institute of Advanced Industrial Science and Technology (AIST), Japan
Shohei Nobuhara	Kyoto University, Japan
Hiroshi Kawasaki	Kyushu University, Japan

W04 – Workshop on Visual Learning and Embodied Agents in Simulation Environments

Peter Anderson	Georgia Institute of Technology, USA
Manolis Savva	Facebook AI Research and Simon Fraser University, USA
Angel X. Chang	Eloquent Labs and Simon Fraser University, USA
Saurabh Gupta	University of California, Berkeley, USA
Amir R. Zamir	Stanford University and University of California, Berkeley, USA
Stefan Lee	Georgia Institute of Technology, USA
Samyak Datta	Georgia Institute of Technology, USA
Li Yi	Stanford University, USA
Hao Su	University of California, San Diego, USA
Qixing Huang	The University of Texas at Austin, USA
Cewu Lu	Shanghai Jiao Tong University, China
Leonidas Guibas	Stanford University, USA

W05 – Bias Estimation in Face Analytics

Rama Chellappa	University of Maryland, USA
Nalini Ratha	IBM Watson Research Center, USA
Rogerio Feris	IBM Watson Research Center, USA
Michele Merler	IBM Watson Research Center, USA
Vishal Patel	Johns Hopkins University, USA

W06 – 4th International Workshop on Recovering 6D Object Pose

Tomas Hodan	Czech Technical University in Prague, Czechia
Rigas Kouskouridas	Scape Technologies, UK
Krzysztof Walas	Poznan University of Technology, Poland
Tae-Kyun Kim	Imperial College London, UK
Jiří Matas	Czech Technical University in Prague, Czechia
Carsten Rother	Heidelberg University, Germany
Frank Michel	Technical University Dresden, Germany
Vincent Lepetit	University of Bordeaux, France
Ales Leonardis	University of Birmingham, UK
Carsten Steger	Technical University of Munich, MVTec, Germany
Caner Sahin	Imperial College London, UK

W07 – Second International Workshop on Computer Vision for UAVs

Kristof Van Beeck	KU Leuven, Belgium
Tinne Tuytelaars	KU Leuven, Belgium
Davide Scaramuzza	ETH Zurich, Switzerland
Toon Goedemé	KU Leuven, Belgium

W08 – 5th Transferring and Adapting Source Knowledge in Computer Vision and Second VisDA Challenge

Tatiana Tommasi	Italian Institute of Technology, Italy
David Vázquez	Element AI, Canada
Kate Saenko	Boston University, USA
Ben Usman	Boston University, USA
Xingchao Peng	Boston University, USA
Judy Hoffman	Facebook AI Research, USA
Neela Kaushik	Boston University, USA
Antonio M. López	Universitat Autònoma de Barcelona and Computer Vision Center, Spain
Wen Li	ETH Zurich, Switzerland
Francesco Orabona	Boston University, USA

W09 – PoseTrack Challenge: Articulated People Tracking in the Wild

Mykhaylo Andriluka	Google Research, Switzerland
Umar Iqbal	University of Bonn, Germany
Anton Milan	Amazon, Germany
Leonid Pishchulin	Max Planck Institute for Informatics, Germany
Christoph Lassner	Amazon, Germany
Eldar Insafutdinov	Max Planck Institute for Informatics, Germany
Siyu Tang	Max Planck Institute for Intelligent Systems, Germany
Juergen Gall	University of Bonn, Germany
Bernt Schiele	Max Planck Institute for Informatics, Germany

W10 – Workshop on Objectionable Content and Misinformation

Cristian Canton Ferrer	Facebook, USA
Matthias Niessner	Technical University of Munich, Germany
Paul Natsev	Google, USA
Marius Vlad	Google, Switzerland

W11 – 9th International Workshop on Human Behavior Understanding

Xavier Alameda-Pineda	Inria Grenoble, France
Elisa Ricci	Fondazione Bruno Kessler and University of Trento, Italy
Albert Ali Salah	Boğaziçi University, Turkey
Nicu Sebe	University of Trento, Italy
Shuicheng Yan	National University of Singapore, Singapore

W12 – First Person in Context Workshop and Challenge

Si Liu	Beihang University, China
Jiashi Feng	National University of Singapore, Singapore
Jizhong Han	Institute of Information Engineering, China
Shuicheng Yan	National University of Singapore, Singapore
Yao Sun	Institute of Information Engineering, China

Yue Liao	Institute of Information Engineering, China
Lejian Ren	Institute of Information Engineering, China
Guanghui Ren	Institute of Information Engineering, China

W13 – 4th Workshop on Computer Vision for Art Analysis

Stuart James	Istituto Italiano di Tecnologia, Italy and University College London, UK
Leonardo Impett	EPFL, Switzerland and Biblioteca Hertziana, Max Planck Institute for Art History, Italy
Peter Hall	University of Bath, UK
João Paulo Costeira	Instituto Superior Tecnico, Portugal
Peter Bell	Friedrich-Alexander-University Nürnberg, Germany
Alessio Del Bue	Istituto Italiano di Tecnologia, Italy

W14 – First Workshop on Fashion, Art, and Design

Hui Wu	IBM Research AI, USA
Negar Rostamzadeh	Element AI, Canada
Leonidas Lefakis	Zalando Research, Germany
Joy Tang	Markable, USA
Rogerio Feris	IBM Research AI, USA
Tamara Berg	UNC Chapel Hill/Shopagon Inc., USA
Luba Elliott	Independent Curator/Researcher/Producer
Aaron Courville	MILA/University of Montreal, Canada
Chris Pal	MILA/PolyMTL, Canada
Sanja Fidler	University of Toronto, Canada
Xavier Snelgrove	Element AI, Canada
David Vazquez	Element AI, Canada
Julia Lasserre	Zalando Research, Germany
Thomas Boquet	Element AI, Canada
Nana Yamazaki	Zalando SE, Germany

W15 – Anticipating Human Behavior

Juergen Gall	University of Bonn, Germany
Jan van Gemert	Delft University of Technology, The Netherlands
Kris Kitani	Carnegie Mellon University, USA

W16 – Third Workshop on Geometry Meets Deep Learning

Xiaowei Zhou	Zhejiang University, China
Emanuele Rodolà	Sapienza University of Rome, Italy
Jonathan Masci	NNAISENSE, Switzerland
Kosta Derpanis	Ryerson University, Canada

W17 – First Workshop on Brain-Driven Computer Vision

Simone Palazzo	University of Catania, Italy
Isaak Kavasidis	University of Catania, Italy
Dimitris Kastaniotis	University of Patras, Greece
Stavros Dimitriadis	Cardiff University, UK

W18 – Second Workshop on 3D Reconstruction Meets Semantics

Radim Tylecek	University of Edinburgh, UK
Torsten Sattler	ETH Zurich, Switzerland
Thomas Brox	University of Freiburg, Germany
Marc Pollefeys	ETH Zurich/Microsoft, Switzerland
Robert B. Fisher	University of Edinburgh, UK
Theo Gevers	University of Amsterdam, Netherlands

W19 – Third International Workshop on Video Segmentation

Pablo Arbelaez	Universidad de los Andes, Columbia
Thomas Brox	University of Freiburg, Germany
Fabio Galasso	OSRAM GmbH, Germany
Iasonas Kokkinos	University College London, UK
Fuxin Li	Oregon State University, USA

W20 – PeopleCap 2018: Capturing and Modeling Human Bodies, Faces, and Hands

Gerard Pons-Moll	MPI for Informatics and Saarland Informatics Campus, Germany
Jonathan Taylor	Google, USA

W21 – Workshop on Shortcomings in Vision and Language

Dhruv Batra	Georgia Institute of Technology and Facebook AI Research, USA
Raffaella Bernardi	University of Trento, Italy
Raquel Fernández	University of Amsterdam, The Netherlands
Spandana Gella	University of Edinburgh, UK
Kushal Kafle	Rochester Institute of Technology, USA
Moin Nabi	SAP SE, Germany
Stefan Lee	Georgia Institute of Technology, USA

W22 – Second YouTube-8M Large-Scale Video Understanding Workshop

Apostol (Paul) Natsev	Google Research, USA
Rahul Sukthankar	Google Research, USA
Joonseok Lee	Google Research, USA
George Toderici	Google Research, USA

W23 – Second International Workshop on Compact and Efficient Feature Representation and Learning in Computer Vision

Jie Qin	ETH Zurich, Switzerland
Li Liu	National University of Defense Technology, China and University of Oulu, Finland
Li Liu	Inception Institute of Artificial Intelligence, UAE
Fan Zhu	Inception Institute of Artificial Intelligence, UAE
Matti Pietikäinen	University of Oulu, Finland
Luc Van Gool	ETH Zurich, Switzerland

W24 – 5th Women in Computer Vision Workshop

Zeynep Akata	University of Amsterdam, The Netherlands
Dena Bazazian	Computer Vision Center, Spain
Yana Hasson	Inria, France
Angjoo Kanazawa	UC Berkeley, USA
Hildegard Kuehne	University of Bonn, Germany
Gül Varol	Inria, France

W25 – Perceptual Image Restoration and Manipulation Workshop and Challenge

Yochai Blau	Technion – Israel Institute of Technology, Israel
Roey Mechrez	Technion – Israel Institute of Technology, Israel
Radu Timofte	ETH Zurich, Switzerland
Tomer Michaeli	Technion – Israel Institute of Technology, Israel
Lihi Zelnik-Manor	Technion – Israel Institute of Technology, Israel

W26 – Egocentric Perception, Interaction, and Computing

Dima Damen	University of Bristol, UK
Giuseppe Serra	University of Udine, Italy
David Crandall	Indiana University, USA
Giovanni Maria Farinella	University of Catania, Italy
Antonino Furnari	University of Catania, Italy

W27 – Vision Meets Drone: A Challenge

Pengfei Zhu	Tianjin University, China
Longyin Wen	JD Finance, USA
Xiao Bian	GE Global Research, USA
Haibin Ling	Temple University, USA

W28 – 11th Perceptual Organization in Computer Vision Workshop on Action, Perception, and Organization

Deepak Pathak	UC Berkeley, USA
Bharath Hariharan	Cornell University, USA

W29 – AutoNUE: Autonomous Navigation in Unconstrained Environments

Manmohan Chandraker	University of California San Diego, USA
C. V. Jawahar	IIIT Hyderabad, India
Anoop M. Namboodiri	IIIT Hyderabad, India
Srikumar Ramalingam	University of Utah, USA
Anbumani Subramanian	Intel, Bangalore, India

W30 – ApolloScape: Vision-Based Navigation for Autonomous Driving

Peng Wang	Baidu Research, USA
Ruigang Yang	Baidu Research, China
Andreas Geiger	ETH Zurich, Switzerland
Hongdong Li	Australian National University, Australia
Alan Yuille	The Johns Hopkins University, USA

W31 – 6th International Workshop on Assistive Computer Vision and Robotics

Giovanni Maria Farinella	University of Catania, Italy
Marco Leo	National Research Council of Italy, Italy
Gerard G. Medioni	University of Southern California, USA
Mohan Trivedi	University of California, USA

W32 – 4th International Workshop on Observing and Understanding Hands in Action

Iason Oikonomidis	Foundation for Research and Technology, Greece
Guillermo Garcia-Hernando	Imperial College London, UK
Angela Yao	National University of Singapore, Singapore
Antonis Argyros	University of Crete/Foundation for Research and Technology, Greece
Vincent Lepetit	University of Bordeaux, France
Tae-Kyun Kim	Imperial College London, UK

W33 – Bioimage Computing

Jens Rittscher	University of Oxford, UK
Anna Kreshuk	University of Heidelberg, Germany
Florian Jug	Max Planck Institute CBG, Germany

W34 – First Workshop on Interactive and Adaptive Learning in an Open World

Erik Rodner	Carl Zeiss AG, Germany
Alexander Freytag	Carl Zeiss AG, Germany
Vittorio Ferrari	Google, Switzerland/University of Edinburgh, UK
Mario Fritz	CISPA Helmholtz Center i.G., Germany
Uwe Franke	Daimler AG, Germany
Terrence Boult	University of Colorado, Colorado Springs, USA

Juergen Gall	University of Bonn, Germany
Walter Scheirer	University of Notre Dame, USA
Angela Yao	University of Bonn, Germany

W35 – First Multimodal Learning and Applications Workshop

Paolo Rota	University of Trento, Italy
Vittorio Murino	Istituto Italiano di Tecnologia, Italy
Michael Yang	University of Twente, The Netherlands
Bodo Rosenhahn	Leibniz-Universität Hannover, Germany

W36 – What Is Optical Flow for?

Fatma Güney	Oxford University, UK
Laura Sevilla-Lara	Facebook Research, USA
Deqing Sun	NVIDIA, USA
Jonas Wulff	Massachusetts Institute of Technology, USA

W37 – Vision for XR

Richard Newcombe	Facebook Reality Labs, USA
Chris Sweeney	Facebook Reality Labs, USA
Julian Straub	Facebook Reality Labs, USA
Jakob Engel	Facebook Reality Labs, USA
Michael Goesele	Technische Universität Darmstadt, Germany

W38 – Open Images Challenge Workshop

Vittorio Ferrari	Google AI, Switzerland
Alina Kuznetsova	Google AI, Switzerland
Jordi Pont-Tuset	Google AI, Switzerland
Matteo Malloci	Google AI, Switzerland
Jasper Uijlings	Google AI, Switzerland
Jake Walker	Google AI, Switzerland
Rodrigo Benenson	Google AI, Switzerland

W39 – VizWiz Grand Challenge: Answering Visual Questions from Blind People

Danna Gurari	University of Texas at Austin, USA
Kristen Grauman	University of Texas at Austin, USA
Jeffrey P. Bigham	Carnegie Mellon University, USA

W40 – 360° Perception and Interaction

Min Sun	National Tsing Hua University, Taiwan
Yu-Chuan Su	University of Texas at Austin, USA
Wei-Sheng Lai	University of California, Merced, USA
Liwei Chan	National Chiao Tung University, USA
Hou-Ning Hu	National Tsing Hua University, Taiwan
Silvio Savarese	Stanford University, USA

Kristen Grauman	University of Texas at Austin, USA
Ming-Hsuan Yang	University of California, Merced, USA

W41 – Joint COCO and Mapillary Recognition Challenge Workshop

Tsung-Yi Lin	Google Brain, USA
Genevieve Patterson	Microsoft Research, USA
Matteo R. Ronchi	Caltech, USA
Yin Cui	Cornell, USA
Piotr Dollár	Facebook AI Research, USA
Michael Maire	TTI-Chicago, USA
Serge Belongie	Cornell, USA
Lubomir Bourdev	WaveOne, Inc., USA
Ross Girshick	Facebook AI Research, USA
James Hays	Georgia Tech, USA
Pietro Perona	Caltech, USA
Deva Ramanan	CMU, USA
Larry Zitnick	Facebook AI Research, USA
Riza Alp Guler	Inria, France
Natalia Neverova	Facebook AI Research, France
Vasil Khalidov	Facebook AI Research, France
Iasonas Kokkinos	Facebook AI Research, France
Samuel Rota Bulò	Mapillary Research, Austria
Lorenzo Porzi	Mapillary Research, Austria
Peter Kontschieder	Mapillary Research, Austria
Alexander Kirillov	Heidelberg University, Germany
Holger Caesar	University of Edinburgh, UK
Jasper Uijlings	Google Research, UK
Vittorio Ferrari	University of Edinburgh and Google Research, UK

W42 – First Large-Scale Video Object Segmentation Challenge

Ning Xu	Adobe Research, USA
Linjie Yang	SNAP Research, USA
Yuchen Fan	University of Illinois at Urbana-Champaign, USA
Jianchao Yang	SNAP Research, USA
Weiyao Lin	Shanghai Jiao Tong University, China
Michael Ying Yang	University of Twente, The Netherlands
Brian Price	Adobe Research, USA
Jiebo Luo	University of Rochester, USA
Thomas Huang	University of Illinois at Urbana-Champaign, USA

W43 – WIDER Face and Pedestrian Challenge

Chen Change Loy	Nanyang Technological University, Singapore
Dahua Lin	The Chinese University of Hong Kong, SAR China
Wanli Ouyang	University of Sydney, Australia
Yuanjun Xiong	Amazon Rekognition, USA
Shuo Yang	Amazon Rekognition, USA
Qingqiu Huang	The Chinese University of Hong Kong, SAR China
Dongzhan Zhou	SenseTime, China
Wei Xia	Amazon Rekognition, USA
Quanquan Li	SenseTime, China
Ping Luo	The Chinese University of Hong Kong, SAR China
Junjie Yan	SenseTime, China

Contents – Part IV

W19 – Third International Workshop on Video Segmentation

Fast Semantic Segmentation on Video Using Block Motion-Based Feature Interpolation 3
Samvit Jain and Joseph E. Gonzalez

Video Object Segmentation with Referring Expressions 7
Anna Khoreva, Anna Rohrbach, and Bernt Schiele

W20 – PeopleCap 2018: Capturing and Modeling Human Bodies, Faces and Hands

MobileFace: 3D Face Reconstruction with Efficient CNN Regression 15
Nikolai Chinaev, Alexander Chigorin, and Ivan Laptev

A Kinematic Chain Space for Monocular Motion Capture 31
Bastian Wandt, Hanno Ackermann, and Bodo Rosenhahn

Non-rigid 3D Shape Registration Using an Adaptive Template 48
Hang Dai, Nick Pears, and William Smith

3D Human Body Reconstruction from a Single Image via Volumetric Regression 64
Aaron S. Jackson, Chris Manafas, and Georgios Tzimiropoulos

Can 3D Pose Be Learned from 2D Projections Alone? 78
Dylan Drover, Rohith M. V, Ching-Hang Chen, Amit Agrawal, Ambrish Tyagi, and Cong Phuoc Huynh

W21 – Shortcomings in Vision and Language

Towards a Fair Evaluation of Zero-Shot Action Recognition Using External Data 97
Alina Roitberg, Manuel Martinez, Monica Haurilet, and Rainer Stiefelhagen

MoQA – A Multi-modal Question Answering Architecture 106
Monica Haurilet, Ziad Al-Halah, and Rainer Stiefelhagen

Pre-gen Metrics: Predicting Caption Quality Metrics Without Generating Captions 114
Marc Tanti, Albert Gatt, and Adrian Muscat

Quantifying the Amount of Visual Information Used by Neural Caption Generators 124
Marc Tanti, Albert Gatt, and Kenneth P. Camilleri

Distinctive-Attribute Extraction for Image Captioning 133
Boeun Kim, Young Han Lee, Hyedong Jung, and Choongsang Cho

Knowing Where to Look? Analysis on Attention of Visual Question Answering System 145
Wei Li, Zehuan Yuan, Xiangzhong Fang, and Changhu Wang

Knowing When to Look for What and Where: Evaluating Generation of Spatial Descriptions with Adaptive Attention 153
Mehdi Ghanimifard and Simon Dobnik

How Clever Is the FiLM Model, and How Clever Can it Be?. 162
Alexander Kuhnle, Huiyuan Xie, and Ann Copestake

Image-Sensitive Language Modeling for Automatic Speech Recognition 173
Kata Naszádi, Youssef Oualil, and Dietrich Klakow

Adding Object Detection Skills to Visual Dialogue Agents. 180
Gabriele Bani, Davide Belli, Gautier Dagan, Alexander Geenen, Andrii Skliar, Aashish Venkatesh, Tim Baumgärtner, Elia Bruni, and Raquel Fernández

W22 – 2nd YouTube-8M Large-Scale Video Understanding Workshop

The 2nd YouTube-8M Large-Scale Video Understanding Challenge 193
Joonseok Lee, Apostol (Paul) Natsev, Walter Reade, Rahul Sukthankar, and George Toderici

NeXtVLAD: An Efficient Neural Network to Aggregate Frame-Level Features for Large-Scale Video Classification 206
Rongcheng Lin, Jing Xiao, and Jianping Fan

Non-local NetVLAD Encoding for Video Classification. 219
Yongyi Tang, Xing Zhang, Jingwen Wang, Shaoxiang Chen, Lin Ma, and Yu-Gang Jiang

Learnable Pooling Methods for Video Classification 229
Sebastian Kmiec, Juhan Bae, and Ruijian An

Constrained-Size Tensorflow Models for YouTube-8M Video Understanding Challenge 239
Tianqi Liu and Bo Liu

Label Denoising with Large Ensembles of Heterogeneous Neural Networks 250
Pavel Ostyakov, Elizaveta Logacheva, Roman Suvorov, Vladimir Aliev, Gleb Sterkin, Oleg Khomenko, and Sergey I. Nikolenko

Hierarchical Video Frame Sequence Representation with Deep Convolutional Graph Network 262
Feng Mao, Xiang Wu, Hui Xue, and Rong Zhang

Training Compact Deep Learning Models for Video Classification Using Circulant Matrices 271
Alexandre Araujo, Benjamin Negrevergne, Yann Chevaleyre, and Jamal Atif

Towards Good Practices for Multi-modal Fusion in Large-Scale Video Classification 287
Jinlai Liu, Zehuan Yuan, and Changhu Wang

Building A Size Constrained Predictive Models for Video Classification 297
Miha Skalic and David Austin

Temporal Attention Mechanism with Conditional Inference for Large-Scale Multi-label Video Classification 306
Eun-Sol Kim, Kyoung-Woon On, Jongseok Kim, Yu-Jung Heo, Seong-Ho Choi, Hyun-Dong Lee, and Byoung-Tak Zhang

Approach for Video Classification with Multi-label on YouTube-8M Dataset 317
Kwangsoo Shin, Junhyeong Jeon, Seungbin Lee, Boyoung Lim, Minsoo Jeong, and Jongho Nang

Learning Video Features for Multi-label Classification 325
Shivam Garg

Large-Scale Video Classification with Feature Space Augmentation Coupled with Learned Label Relations and Ensembling 338
Choongyeun Cho, Benjamin Antin, Sanchit Arora, Shwan Ashrafi, Peilin Duan, Dang The Huynh, Lee James, Hang Tuan Nguyen, Mojtaba Solgi, and Cuong Van Than

W23 – 2nd International Workshop on Compact and Efficient Feature Representation and Learning in Computer Vision

Multi-style Generative Network for Real-Time Transfer 349
Hang Zhang and Kristin Dana

Frustratingly Easy Trade-off Optimization Between Single-Stage and Two-Stage Deep Object Detectors 366
Petru Soviany and Radu Tudor Ionescu

Targeted Kernel Networks: Faster Convolutions with Attentive Regularization 379
Kashyap Chitta

Small Defect Detection Using Convolutional Neural Network Features and Random Forests 398
Xinghui Dong, Chris J. Taylor, and Tim F. Cootes

Compact Deep Aggregation for Set Retrieval 413
Yujie Zhong, Relja Arandjelović, and Andrew Zisserman

Adversarial Network Compression. 431
Vasileios Belagiannis, Azade Farshad, and Fabio Galasso

Target Aware Network Adaptation for Efficient Representation Learning 450
Yang Zhong, Vladimir Li, Ryuzo Okada, and Atsuto Maki

Learning CCA Representations for Misaligned Data 468
Hichem Sahbi

Learning Relationship-Aware Visual Features 486
Nicola Messina, Giuseppe Amato, Fabio Carrara, Fabrizio Falchi, and Claudio Gennaro

DNN Feature Map Compression Using Learned Representation over GF(2) 502
Denis Gudovskiy, Alec Hodgkinson, and Luca Rigazio

LBP-Motivated Colour Texture Classification 517
Raquel Bello-Cerezo, Paul Fieguth, and Francesco Bianconi

Discriminative Feature Selection by Optimal Manifold Search for Neoplastic Image Recognition 534
Hayato Itoh, Yuichi Mori, Masashi Misawa, Masahiro Oda, Shin-Ei Kudo, and Kensaku Mori

Fast, Visual and Interactive Semi-supervised Dimensionality Reduction 550
Dimitris Spathis, Nikolaos Passalis, and Anastasios Tefas

Efficient Texture Retrieval Using Multiscale Local Extrema Descriptors and Covariance Embedding 564
Minh-Tan Pham

Extended Non-local Feature for Visual Saliency Detection in Low Contrast Images 580
Xin Xu and Jie Wang

Incomplete Multi-view Clustering via Graph Regularized Matrix Factorization 593
Jie Wen, Zheng Zhang, Yong Xu, and Zuofeng Zhong

GA-Based Filter Selection for Representation in Convolutional Neural Networks 609
Junbong Kim, Minki Lee, Jongeun Choi, and Kisung Seo

Active Descriptor Learning for Feature Matching 619
Aziz Koçanaoğulları and Esra Ataer-Cansızoğlu

A Joint Generative Model for Zero-Shot Learning 631
Rui Gao, Xingsong Hou, Jie Qin, Li Liu, Fan Zhu, and Zhao Zhang

W24 – Women in Computer Vision Workshop

WiCV at ECCV2018: The Fifth Women in Computer Vision Workshop 649
Zeynep Akata, Dena Bazazian, Yana Hasson, Angjoo Kanazawa, Hildegard Kuehne, and Gül Varol

Gait Energy Image Reconstruction from Degraded Gait Cycle Using Deep Learning 654
Maryam Babaee, Linwei Li, and Gerhard Rigoll

Hierarchical Video Understanding 659
Farzaneh Mahdisoltani, Roland Memisevic, and David Fleet

Fine-Grained Vehicle Classification with Unsupervised Parts Co-occurrence Learning 664
Sara Elkerdawy, Nilanjan Ray, and Hong Zhang

Multiple Wavelet Pooling for CNNs 671
Aina Ferrà, Eduardo Aguilar, and Petia Radeva

Automated Facial Wrinkles Annotator 676
Moi Hoon Yap, Jhan Alarifi, Choon-Ching Ng, Nazre Batool, and Kevin Walker

Deep Learning of Appearance Models for Online Object Tracking 681
Mengyao Zhai, Lei Chen, Greg Mori, and Mehrsan Javan Roshtkhari

Towards Cycle-Consistent Models for Text and Image Retrieval 687
Marcella Cornia, Lorenzo Baraldi, Hamed R. Tavakoli, and Rita Cucchiara

From Attribute-Labels to Faces: Face Generation Using a Conditional Generative Adversarial Network 692
Yaohui Wang, Antitza Dantcheva, and Francois Bremond

Optimizing Body Region Classification with Deep Convolutional Activation Features 699
Obioma Pelka, Felix Nensa, and Christoph M. Friedrich

Efficient Interactive Multi-object Segmentation in Medical Images 705
Leissi Margarita Castañeda Leon and Paulo André Vechiatto de Miranda

Cross-modal Embeddings for Video and Audio Retrieval 711
Didac Surís, Amanda Duarte, Amaia Salvador, Jordi Torres, and Xavier Giró-i-Nieto

Understanding Center Loss Based Network for Image Retrieval with Few Training Data 717
Pallabi Ghosh and Larry S. Davis

End-to-End Trained CNN Encoder-Decoder Networks for Image Steganography 723
Atique ur Rehman, Rafia Rahim, Shahroz Nadeem, and Sibt ul Hussain

Cancelable Knuckle Template Generation Based on LBP-CNN 730
Avantika Singh, Shreya Hasmukh Patel, and Aditya Nigam

A 2.5D Deep Learning-Based Approach for Prostate Cancer Detection on T2-Weighted Magnetic Resonance Imaging 734
Ruba Alkadi, Ayman El-Baz, Fatma Taher, and Naoufel Werghi

GreenWarps: A Two-Stage Warping Model for Stitching Images Using Diffeomorphic Meshes and Green Coordinates 740
Geethu Miriam Jacob and Sukhendu Das

Author Index 745