Dr Andrea Vedaldi

Department of Engineering Science
University of Oxford
Parks Road, Oxford, UK
OX1 3PJ

E-mail vedaldi@robots.ox.ac.uk
WWW http://www.robots.ox.ac.uk/~vedaldi
Tel. +44 1865 273 127

Research interests

Image and video understanding, unsupervised machine learning, unsupervised 3D geometry, motion and physics.

Professional experience

2023– **Research Scientist**, at *Meta AI*, London, UK.

2018–23 **Research Scientist**, at *Facebook AI Research* (FAIR), London, UK.

2020– **Professor of Computer Vision and Machine Learning** at the Engineering Science department, *University of Oxford*, UK.

2012–20 **Associate Professor** in Engineering Science at the *University of Oxford*, UK.

2012–19 **Tutorial Fellow** at *New College*, Oxford, UK.

2010–12 **Junior Research Fellow** at the *University of Oxford*, UK.

2008–10 **Postdoctoral Research Assistant** at the *University of Oxford*, UK.

Education

2005–08 **Doctor of Philosophy (PhD)**
University of California at Los Angeles
Department of Computer Science, Los Angeles, California
Thesis: "Invariant representations and learning for computer vision"
Committee: Alan Yuille, Song-Chun Zhu, Luiminita Vese, Serge Belongie, Stefano Soatto (chair)

2003–05 **Master of science (MSc) in computer science**
University of California at Los Angeles
Department of Computer Science, Los Angeles, California
GPA 4/4, Advisor: Stefano Soatto

1998–03 **Bachelor of science (laurea quinquennale) in information engineering**
University of Padua
Department of Information Engineering, Padua, Italy
Thesis: "Modelli stocastici per il riconoscimento di sistemi di gesti complessi"
Highest honours (summa cum laude), Advisor: Ruggero Frezza

Grants

2022–26 €2,311,847, PI, "Unsupervised Perception", ERC Consolidator Grant.

2020–21 $80,000, PI, "Large-Scale Understanding of Self-Supervised Image Feature Representation Learning", Amazon Research Award (research led by Yuki Asano and Christian Rupprecht).

2020–25 £5,912,097, Co-I, EPSRC programme grant "Visual AI".

2019–22 £877,887, PI, "Oxford-Continental II", Continental Corp.

2018–19 £200,000, Co-I, "Face recognition", Toshiba.

2019–21 £200,000, Co-I, "Fine-grained detection of novel objects in cluttered environments", Nielsen.

2018–22 £860,000, PI, Facebook DPhil scholarships.

2018–19 $150,000, PI, "Productive Artificial Intelligence", Amazon Research Award.

2016–20 £61,000, Co-I, "Environmental Bayesian Optimization of General Purpose Neural Networks", MathWorks.

2016–20 £720,592, PI, "Oxford-Continental I", Continental Corp.

2015–20 £4,466,184, Co-I, EPSRC programme grant "Seebibyte".

2015–20 €1,497,271, PI, "Integrated and Detailed Image Understanding", ERC Starting Grant.

2014–15 £100,000, "Detailed and Deep Image Understanding", EPSRC First Grant.

2015–17 £90,000, PI, "Fine-grained object detection and description for scene understanding", Xerox Foundation/Naver PhD Sponsorship.

2014–17 $300,000, Co-I, "Selecting, Stitching & Identifying Digital Data Used In Integrity Management", BP.

2014–16 £90,000, PI, "Structured Regression from Dynamical Visual Textures", BP.

2012–14 $90,000, Co-I, "Comprehensive and Efficient Recognition of Visually Similar Categories in Images", Xerox Foundation UAC gift.

2012–13 €21,420, PI, PASCAL Harvest Grant for VLFeat.

Awards and other recognitions

2021 Best student paper, *British Machine Vision Conference*.

2020 SIGMM Test of time award, *ACM MULTIMEDIA*.

2020 Best paper award, *Conference on Computer Vision and Pattern Recognition* (CVPR).

2016 Best workshop paper, *ECCV Workshop on Geometry Meets Deep Learning*.

2015 IEEE PAMI Mark Everingham Prize, *ICCV*.

2014 Best paper award, *British Machine Vision Conference* (BMVC).

2014 Outstanding Reviewer Award, IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

2013 Best Reviewer Award, Conf. on Neural Information Processing Systems (NeurIPS).

2013 Best Reviewer Award, International Conference on Computer Vision (ICCV).

2012 Outstanding Reviewer Award, IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Outstanding Reviewer Award, IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

2011 W. W. Spooner Junior Research Fellow, New College, Oxford.

2011 Best Poster Honorable Mention, *British Machine Vision Conference* (BMVC).

2010 Oxford Violette and Samuel Glasstone Research Fellowships in Science.

2010 Winner of the "ACM Multimedia Open Source Software Competition" for VLFeat, with Brian Fulkerson, at the ACM international conference on Multimedia.

2009 Winner of the "PASCAL Visual Object Classification Challenge" for object detection, with V. Gulshan, M. Varma, and A. Zisserman, at the International Conference on Computer Vision (ICCV).

2008 "Outstanding Doctor of Philosophy in Computer Science", *The Henry Samueli School of Engineering and Applied Science, University of California at Los Angeles*.

2005 "Outstanding Master of Science in Computer Science", *The Henry Samueli School of Engineering and Applied Science, University of California at Los Angeles*.

Service to the academic community

2026 **General Chair.** European Conference on Computer Vision (ECCV).

2020 **Program Chair.** European Conference on Computer Vision (ECCV).

- **Area Chair.** International Conference on Computer Vision (ICCV), 2011, British Machine Vision Conference (BMVC), 2012, 2013, 2017, Computer Vision and Pattern Recognition (CVPR) 2013, 2015, European Conference on Computer Vision (ECCV) 2014, 2016, 2018.

2017 **Local Chair.** International Conference on Computer Vision (ICCV).

2018 **Tutorial Chair.** Computer Vision and Pattern Recognition (CVPR).

2019 **Workshop Chair.** Computer Vision and Pattern Recognition (CVPR).

- **Associate editor.** Associate editor of the IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2014–19. CVIU special issue on Deep Learning, 2016.

- **Reviewer for international journals.** IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), International Journal on Computer Vision (IJCV), IEEE Transactions on Image Processing (TIP), Computer Vision and Image Understanding (CVIU), Neurocomputing.

- **Reviewer for international conferences.** IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2005–2012, 2014–17 European Conference on Computer Vision (ECCV) 2006, 2008, 2010, 2012, 2016, International Conference on Computer Vision (ICCV) 2005, 2007, 2009, 2013, 2015, International Conference on Neural Information Processing Systems (NeurIPS) 2007–14, International Conference on Robotics Science and Systems (ICRSS) 2012, International Conference on Learning Representations (ICLR) 2016.

- **Reviewer for national conferences.** Indian Conference on Computer Vision, Graphics & Image Processing (ICCVGP).

- **Reviewer for funding schemes.** EPSRC Standard Proposal, 2015, 2016, 2017, ERC First and Consolidators Grants, 2016, 2017, 2019, Leverhulme Trust, 2013, CHIST-ERA, 2012, Vienna Science and Technology Fund (WWTF), 2011, EPSRC First Grant, 2014.

Software projects

I am the leading author of the VLFeat library of computer vision algorithms (http://www.vlfeat.org/). The library is a popular tool for fast prototyping in computer vision research. For this project, I was awarded the PAMI Mark Everingham Prize in 2015.

I am also the leading author of the MatConvNet library (http://www.vlfeat.org/matconvnet), a MATLAB/C/CUDA library for fast prototyping with Convolutional Neural Networks.

Event organisation

I was **team leader** for the Johns Hopkins CLSP Summer Research Workshop "Towards a Detailed Understanding of Objects and Scenes in Natural Images", June 11 — August 7 2012, in collaboration with Subhransu Maji, Matthew Blaschko, Iasonas Kokkinos, Ben Taskar.

International workshops

I have co-organised the following international workshops:

2022 "Neural Geometry and Rendering: Advances and the Common Objects in 3D Challenge"at the *European Conference on Computer Vision* (ECCV), with David Novotny, Shangzhe Wu, Roman Shapovalov, Samarth Sinha, Natalia Neverova and Jitendra Malik. https://ngr-co3d.github.io

2022 "AV4D: Visual Learning of Sounds in Spaces" at the *European Conference on Computer Vision* (ECCV), with Changan Chen, Ruohan Gano, Andrew Owens, David Harwath, Chuang Gan, Antonio Torralba and Kristen Grauman. https://av4d.org/eccv22

2021 "Workshop on Autonomous Driving" at the *Computer Vision and Pattern Recognition* (CVPR), with Yuning Chai, Henrik Kretzschmar, Yang Song, Lukas Neumann, Andreas Geiger, Dragomir Anguelov, Alexander Liniger, Jose M Alvarez, Fisher Yu, David Vazquez, Antonio M. Lopez, Tomas Pajdla, Luc Van Gool, John Leonard. https://cvpr2021.wad.vision

2020 "Workshop on Adversarial Robustness in the Real World" at the *International Conference on Computer Vision* (ICCV), with Adam Kortylewski, Cihang Xie, Song Bai, Zhaowei Cai, Yingwei Li, Andrei Barbu, Wieland Brendel, Philip H.S. Torr, Rama Chellappa, Alan L. Yuille. https://eccv22-arow.github.io

2020 "Self Supervised Learning: What is Next?" workshop at the *European Conference on Computer Vision* (ECCV), with Christian Rupprecht, Yuki M Asano, Armand Joulin. https://sslwin.org

2020 "Scalability in Autonomous Driving" at the *Computer Vision and Pattern Recognition* (CVPR) with Yuning Chai, Henrik Kretzschmar, Yin Zhou, Pei Sun, Lukas Neumann, Andrea Vedaldi, Andreas Geiger, Dragomir Anguelov. https://sites.google.com/view/cvpr20-scalability

2019 "Neural Architects Workshop" at the *International Conference on Computer Vision* (ICCV), with Samuel Albanie, Li Shen, Jie Hu, Barret Zoph, Andrea Vedaldi, Andrew Zisserman. https://neuralarchitects.org

2018 "Rank Prize Symposium: Geometry and Uncertainty in Deep Learning for Computer Vision", Wordsworth, UK, with R. Cipolla and A. Kendall. https://www.rankprize.org/symposia/optoelectronics/2010-2019/

2017 "Interpreting, Explaining and Visualizing Deep Learning", workshop at *Neural Information Processing System* (NeurIPS), with Klaus-Robert Müller, Lars Kai Hansen, Wojciech Samek and Grégoire Montavon. http://www.interpretable-ml.org/nips2017workshop/

2017 "PASCAL in detail", workshop at the conference on *Computer Vision and Pattern Recognition* (CVPR), with Sanja Fidler, Iasonas Kokkinos, Roozbeh Mottaghi, George Papandreou, Raquel Urtasun, and Alan L. Yuille. https://sites.google.com/view/pasd

2016 "Local Features: State Of The Art, Open Problems And Performance Evaluation", workshop at the *European Conference on Computer Vision* (ECCV), with Jiri Matas, Krystian Mikolajczyk, Tinne Tuytelaars, Vassileios Balntas and Karel Lenc. http://icvl.ee.ic.ac.uk/DescrWorkshop/

Tutorials at international conferences

I have co-organised the following tutorials at international conferences:

2019 "Interpretable Machine Learning for Computer Vision" at the *International Conference on Computer Vision* (ICCV), with Bolei Zhou, Zeynep Akata, Trevor Darrel and Alan L. Yuille.

2018 "Interpretable Machine Learning for Computer Vision" at the conference on *Computer Vision and Pattern Recognition* (CVPR), with Bolei Zhou, Laurens van der Maaten and Been Kim.

2014 "Image Representations, from Shallow to Deep" at the *British Machine Vision Conference* (BMVC).

2014 "Tutorial on Large-Scale Visual Recognition" at the conference on *Computer Vision and Pattern Recognition* (CVPR), with O. Chum, Z. Harchaoui, H. Jegou, F. Perronnin, and S. Lazebnik.

2013 "Tutorial on Large-Scale Visual Recognition" at the conference on *Computer Vision and Pattern Recognition* (CVPR), with O. Chum, Z. Harchaoui, H. Jegou, F. Perronnin, and M. Ranzato.

2012 "Additive Kernels and Explicit Embeddings for Large Scale Computer Vision Problems" at the *European Conference on Computer Vision* (ECCV), with J. Wu, S. Maji, and F. Perronnin.

2012 "Modern features: advances, applications, and software" at the *European Conference on Computer Vision* (ECCV), with J. Matas, K. Mikolajczyk, T. Tuytelaars, C. Schmid, and A. Zisserman.

2010 "VLFeat: An Open and Portable Library of Computer Vision Algorithms" at the *European Conference on Computer Vision* (ECCV), with Brian Fulkerson.

2010 "VLFeat: An Open and Portable Library of Computer Vision Algorithms" at the conference on *Computer Vision and Pattern Recognition* (CVPR), with Brian Fulkerson.

International computer vision schools

I have contributed the following international schools:

2022 Lecture on "Learning 3D Geometry" at the University of Amsterdam, NL.

2022 "International Computer Vision Summer School", Catania, Italy. https://iplab.dmi.unict.it/icvss2022/

2021 "OxML — Oxford Machine Learning Summer School", Oxford, UK.

2020 "OxML — Oxford Machine Learning Summer School", Oxford, UK.

2018 "Medical Imaging Summer School", Favignana, Italy.

2017 "Summer School on Signal Processing Meets Deep Learning", Capri, Italy.

2016 "Medical Imaging Summer School", Favignana, Italy.

2016 "Deep learning for computer vision applications", iV&L Net Training School, Malta.

2015 "Image Representations, from Shallow to Deep", International Computer Vision Summer School, Catania, Italy.

2014 "Vision and Sports Summer Schools", Prague.

2013 "International Computer Vision Summer School", Catania, Italy.

2013 Co-organizer with A. Zisserman of the 4^{th} PAVIS summer school, IIT, Genoa, Italy, on "Large Scale Visual Recognition of Object Instances and Categories".

2012 "Visual recognition of object instances and categories", NAACL summer school, *Johns Hopkins University*, Baltimore, US.

Online courses

I am the leading author of the *VGG Computer vision practicals*[1], which ware widely used in international schools on computer vision.

Media appearances

2020 Digital Trends. https://www.digitaltrends.com/features/facebook-ai-image-recognition/

2014 BBC. http://www.bbc.co.uk/informationandarchives/archivenews/2014/face-recognition-and-new-ways-to-search-for-archive.html

Mentoring and student supervision

Postdoctoral researchers

2023– *Edgar Sucar*.

2022– *Chuanxia Zheng*.

2020– *Eldar Insafutdinov*.

2020– *Iro Laina*.

2020–2022 *Dylan Campbell*, then Professor at Australian National University.

2018–2021 *Han Peng*, then **Professor at Hong Kong University of Sciences and Technology**.

2018–2021 *Dan Xu*.

2018–2021 *Kai Han*, then **Professor at the University of Bristol**.

---

[1]https://sites.google.com/site/vggpracticals/

2018–2021 *Christian Rupprecht*, then **Professor at the University of Oxford**.

2017–2020 *Lukas Neumann*.

2017–2019 *Fatma Guney*, then **Professor at Koç University**.

2016–2019 *Joao F. Henriques*, then **DL at the University of Oxford** supported by a **RAEng fellowship**.

2017–2018 *Maria Klodt*.

2015–2017 *Hakan Bilen*, then **Professor at the University of Edinburgh**.

PhD

2023– *Zihang Lai*, University of Oxford.

2023– *Ruining Li*, University of Oxford.

2023– *Paul Engstler*, University of Oxford.

2022– *Stanislaw Szymanowicz*, University of Oxford, co-supervised with Christian Rupprecht.

2022– *Minghao Chen*, University of Oxford, co-supervised with Iro Laina.

2021– *Laurynas Karazija*, University of Oxford, co-supervised with Christian Rupprecht and Iro Laina.

2021– *Aleksandar Shtedritski*, University of Oxford, co-supervised with Christian Rupprecht.

2020– *Vadim Tschernezki*, University of Oxford, co-supervised with Diane Larlus and Iro Laina.

2020– *Luke Melas-Kyriaz*, University of Oxford, co-supervised with Christian Rupprecht and Iro Laina.

2019–2021 *Mandela Patrik*, University of Oxford, co-supervised with Joao Henriques.

2019– *Subhabrata Choudhury*, University of Oxford.

2019– *Sagar Vaze*, University of Oxford, co-supervised with Andrew Zisserman.

2018–2022 *Shangzhe Wu*, University of Oxford.

2018–2022 *Robert McCraith*, University of Oxford.

2018–2021 *Yuki Asano*, University of Oxford, then **Professor at the University of Amsterdam**.

2018–2021 *Honglie Chen*, University of Oxford.

2017–2022 *Tomas Jakab*, University of Oxford.

2017–2021 *Xu Ji*, University of Oxford, then at **MILA**.

2017–2021 *Oliver Groth*, University of Oxford, co-supervised with Ingmar Posner, then at **Google DeepMind**.

2016–2020 *Sylvestre Rebuffi*, University of Oxford, then at **Google DeepMind**.

2016–2020 *Sebastien Ehrhardt*, University of Oxford.

2016–2020 *Ruth Fong*, University of Oxford, then **faculty at Princeton**.

2016–19 *Dmitry Ulyanov*, Skoltech, then at **Samsung Research**, co-supervised with Victor Lempitsky.

2015–19 *Sam Albanie*, University of Oxford, then **Professor at the University of Cambridge**.

2015–19 *James Thewlis*, University of Oxford.

2015–19 *Ankush Gupta*, University of Oxford, then at **Google DeepMind**, co-supervised with Andrew Zisserman.

2015–18 *David Novotny*, University of Oxford, then at **Facebook AI Research**.

2014–18 *Karel Lenc*, University of Oxford, then at **Google DeepMind**.

2014–18 *Aravindh Mahendran*, University of Oxford, then **Google Brain**.

2012–15 *Max Jaderberg*, University of Oxford, then at **Google DeepMind**, co-supervised with Andrew Zisserman.

2011 *Marco Pedersoli*, University of Barcelona, co-supervised for two terms with Jordi Gonzàlez.

2011–15 *Omkar Parkhi*, University of Oxford, then at **Facebook Applied ML Research**, co-supervised with Andrew Zisserman.

2011–15 *Mircea Cimpoi*, University of Oxford, then at **Google Brain**.

MSc and MEng

I have also supervised undergraduate students in the following **4th year (MSc) project** in Oxford:

2019–20 *Zheyuan Chen*.

2019–20 *Aleksandar Shtedritski*.

2018–19 *Rajeev Shanka*.

2018–19 *Jack Toner*.

2017–18 *DindDing Chen*.

2017–18 *Harry Sargent*.

2016–17 *Natchapol Suebsubanant*.

2016–17 *Jack Frigaard*.

2016–17 *Junni Shou*.

2015–16 *I-Horng Huang*.

2015–16 *Guanjie Wang*.

2015–16 *Zain Khawaja*.

2014–15 *Peter Zengh*.

2014–15 *Matthew Pybus*.

2013–14 *Jai Juneja*.

2013–14 *Tomasz Kaminski*.

2013–14 *Stephen Kyberd*.

2012–13 *Shangqian Lee*.

2012–13 *Asim Jamal*.

2010–11 *John Cant*.

2011–13 *Max Jaderberg*.

I have also supervised the following students:

2009–13 *Mayank Juneja*, MSc, International Institute of Information Technology — Hyderabad.

2008–11 *Omkar Parkhi*, MSc, International Institute of Information Technology — Hyderabad, co-supervised with C. V. Jawahar and Andrew Zisserman.

2010 *Vempati Sreekanth*, MSc, International Institute of Information Technology — Hyderabad, co-supervised with C. V. Jawahar and Andrew Zisserman.

2011–12 *Michal Spisiak*, BMath, 3YP, University of Oxford.

Teaching

Graduate teaching

I teach the following parts of **graduate modules**, which I have created from scratch:

2014– AIMS CDT module on "Learning form Big Data" (six hours), including **practicals**.

2014– AIMS CDT module on "Computer Vision" (six hours), including **practicals**.

Undergraduate teaching

I taught the following **undergraduate modules**, **classes**, and **laboratories**:

2022– New B16 third year module on "Algorithms" (4 lectures) and classes at the *University of Oxford*.

2012–22 New B16 third year module on "Structured programming" (4 lectures) and classes (redesigned from scratch) at the *University of Oxford*.

2012–15 B16 "Structured Programming labs" (redesigned from scratch) at the *University of Oxford*.

2012– C18 fourth year module on "Computer Vision and Deep Learning" (4 lectures) and classes (redesigned from scratch) at the *University of Oxford*.

2012 C19 Machine Learning classes at the *University of Oxford*.

2009–12 B4 Information Engineering laboratories on Computer Vision at the *University of Oxford*.

2010 C4B Computer Vision course, 2010 at the *University of Oxford*.

2005–08 Various lectures and clases on Computer Vision and Computer Graphics at the *University of California at Los Angeles*.

College teaching

I taught at New College as Tutorial Fellow:

2012–15 Full stint (6 hours/week), tutoring the P2 and A2 papers and part of A3.

2015–19 Half stint (3 hours/week) for A2.

2015–19 17 of the 19 4th projects I supervised were part of my college stint.

Examining

2023,24,25 Final examining for the A1 paper.

2016 Final examining for the A2 paper.

2012– B16 and C18 examinations.

2012–19 University admissions as College Tutor for New College.

I have acted as **examiner for the following DPhil/PhD** students:

2024 *Pau de Jorge*, University of Oxford.

2024 *Changan Chen*, University of Texas at Austin.

2024 *Anh Tuan Nguyen*, University of Oxford.

2024 *Ce Liu*, ETH Zurich.

2023 *Christoph Mayer*, ETH Zurich.

2023 *Mariana-Iuliana Georgescu*, University of Bucharest.

2023 *Theo Costain*, University of Oxford.

2022 *Zhao Yang*, University of Oxford.

2022 *Valentin Gabeur*, University of Grenoble.

2022 *Tengda Han*, University of Oxford.

2022 *Nikita Araslanov*, University of Darmstad.

2021 *Marting Engeckle*, University of Oxford.

2021 *Erika Lu*, University of Oxford.

2021 *Dmytro Mishkin*, Czech Technical Institute.

2021 *David Wisth*, University of Oxford.

2020 *Xing Deng*, Imperial College.

2020 *Qizhu Li*, University of Oxford.

2020 *Maxime Berman*, University of Leuven.

2020 *Leonard Berrada*, University of Oxford.

2020 *Ignacio Rocco*, Ecole normale supérieure.

2020 *Chen-Hsuan Lin*, Carnegie Mellon University.

2019 *Rahaf Al Jundi*, University of Leuven.

2019 *Nikita Dvornik*, University of Grenoble.

2019 *Gurkirt Singh*, Oxford Brookes University.

2019 *Anurag Arnab*, University of Oxford.

2018 *Yani Andrew Ioannou*, University of Cambridge.

2018 *Weidi Xie*, University of Oxford.

2018 *Suman Saha*, Brookes University.

2018 *Diane Bouchacourt*, University of Oxford.

2018 *Bohan Zhuang*, Australian National University.

2017 *Sukrit Shankar*, University of Cambridge.

2016 *Saiful Khan*, University of Oxford.

2015 *Yuning Chai*, University of Oxford.

2015 *Tomas Pfister*, University of Oxford.

2015 *Susana Brandão*, Carnegie Mellon University and Técnico Lisboa.

2014 *Zeng Wang*, University of Oxford.

2014 *Tomasz Trzcinski*, EPFL.

2014 *Sadeep Jayasumana*, Australian National University.

2013 *Ziming Zhang*, Oxford Brookes University.

2013 *Yusuf Aytar*, University of Oxford.

2013 *Relja Arandjelović*, University of Oxford, Oxford.

Administration

University administration

I have served in the following **committees** and **panels**:

2015–18 Undergraduate Studies Committee.

2015–17 Information Technology Committee.

2012– Information Engineering Panel.

College administration

As New College Tutor, I have served in the following **committees**:

2012–13 Disciplinary committee.

2013–19 Development committee.

2012–19 Governing Body and Tuition and Research, Warden and Tutors and Admission committees.

Academic development

I have attended the following seminars and courses: "Introduction to Academic Policy and Practice", "Introduction to Teaching at Oxford", "Student Support and Welfare / Academic Life at Oxford: university and college roles and responsibilities", "Research Student Supervision", "The UK Research Environment: funding and conduct of research", "Admissions Interviewing course", "Admissions Interviewing Practice", Oxford, 2012.

Invited talks

2024 "Towards a 3D foundation", keynote at the Int. Conf. on *3D Vision* (3DV). https://3dvconf.github.io/2024/

2023 "A statistical learning perspective on reconstructing the 3D world", talk at the *BrainWorlds Freiburg-Oxford Workshop*. https://brainworlds.uni-freiburg.de

2023 "Understanding egocentric data in 3D" talk at the *3rd International Ego4D Workshop at CVPR*, CA. https://ego4d-data.org/workshops/cvpr23/

2023 "Learning 3D Geometry: From Fusion to Generation" talk at the *CVPR CV4MR Workshop on Computer Vision for Mixed Reality*, CA. https://cv4mr.github.io

2022 "Unsupervised Geometry" talk at the *Rank Prize Workshop*, UK. https://www.rankprize.org/symposia/neural-rendering-in-computer-vision/

2022 "Unsupervised 3D perception" talk at the *NeurIPS Workshop on Self-Supervised Learning — Theory and Practice*, US. https://sslneurips22.github.io

2022 "Unsupervised 3D perception" talk at the *ECCV Workshop on Computer Vision for Metaverse*, IL.

2022 "Unsupervised 3D perception" talk at the *ECCV Workshop on Self Supervised Learning: What is Next?*, IL. https://sslwin.org

2021 "Learning 3D representations" talk at the *ICCV Workshop on Learning 3D Representations for Shape and Appearance*, *International Conference on Computer Vision (ICCV)*, Virtual.

2021 "Discovering actionable interpretations from raw visual data: from 2D clustering to 3D reconstruction" talk at *ELLIS Life / NCT Data Science Seminar*, Heidelberg University, Germany. https://www.dkfz.de/en/datascience/seminar/Vedaldi.html

2021 "Learning 3D objects in the real world" talk at the *CVPR Workshop on Frontiers of Monocular 3D Perception*, Virtual. https://sites.google.com/view/mono3d-workshop

2021 "Learning 3D objects in the real world" talk at the *CVPR Workshop on 3D Scene Understanding for Vision, Graphics and Robotics*, Virtual. https://scene-understanding.com/2021/

2020 "Deep priors" talk at the *Deep Internal Learning: Training with no prior examples* workshop at the European Conference on Computer Vision (ECCV), UK. https://sites.google.com/view/deepinternallearning

2020 "Learning representations and geometry from unlabelled data" talk at the workshop on *Learning form Unlabelled Videos* at the Computer Vision and Pattern Recognition (CVPR), USA. https://sites.google.com/view/luv2020

2019 "Learning semantics and geometry with less supervision" talk at the *Fine-Grained Visual Categorization* workshop at the conference on Computer Vision and Pattern Recognition (CVPR).

2019 "Unsupervised geometry" **keynote** at the *Geometry and Deep Learning* meeting, British Machine Vision Association (BMVA), London, UK.

2017 "Dreaming of electric sheep: how computers see the world", **keynote** at the *Continental Software Conference*, Frankfurt, Germany.

2017 "Visualizing image representations", talk at the *Dagstuhl Seminar*, Dagsthul, Germany.

2017 "Local feature detectors and descriptors in the era of deep learning: practical and theoretical progress', talk at the *CEFRL Workshop on Compact and Efficient Feature Representation and Learning* in conjunction with the *IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, Venice, Italy.

2017 "Universal, unsupervised, and understandable deep image representations", talk at *Naver Research*, Grenoble, France.

2015 "Learning and understanding visual representations", talk at the *Deep Vision Workshop*, conference on Computer Vision and Pattern Recognition (CVPR).

2015 "Learning visual representations", talk at *MPI*, Tübingen, Germany.

2015 "Learning visual representations", talk at *Yandex*, Moscow, Russia.

2014 "The power of learning representations", talk at *Prague Technical University*, CZE.

2014 "Image representations for large-scale visual recognition", talk at *Prague Technical University*, CZE.

2014 "Detailed Image Understanding", talk at *Microsoft Research Redmond*, United States.

2013 "Large Scale Visual Recognition of Object Instances and Categories", talk at the *University of Oulu*, Oulu, Finland.

2013 "Learning to Compare & Compress Local Features and Faces", talk at *Xerox Research Centre Europe*, Grenoble, France.

2011 "Semantic Image Analysis with Structure and Kernels " talk at the *Max Planck Institute*, Saarbrücken, Germany.

2011 "Efficient Structured Modeling for the Interpretation of Images", talk at the *CIS Seminar*, Johns Hopkins University.

2011 "Multiple and fast kernels for the detection of visual object categories", talk at the workshop on *Kernels and Distances for Computer Vision*, International Conference on Computer Vision (ICCV).

2011 "Efficient Structured Modeling for the Interpretation of Images", talk at the *LCSR/ERC Seminar*, Johns Hopkins University.

2010 "Efficient Additive Kernels via Explicit Feature Maps", talk at the *CIS Seminar*, Johns Hopkins University.

2010 "Efficient Additive Kernels via Explicit Feature Maps", talk at the *University of Edinburgh*.

2010 "Efficient Additive Kernels via Explicit Feature Maps", talk at *Heriot-Watt University*.

2010 "A Tale of Two Object Detectors", talk at the *CS seminar*, University of California at San Diego (UCSD).

2009 "Multiple kernels for Object Classification and Detection", talk at the *PASCAL VOC Challenge Workshop*, International Conference on Computer Vision (ICCV).

2009 "Multiple kernels for Object Classification and Detection", talk at Indian Institute of Information Technology (IIIT), Hyderabad.

2009 "The design and implementation of the SIFT feature detector and descriptor", talk at the Indian Institute of Information Technology (IIIT), Hyderabad.

2008 "Invariant Representation and Learning for Computer Vision", talk at the *University of Oxford*.

2008 "Invariant Representation and Learning for Computer Vision", talk at the *École Centrale*.

2007 "Viewpoint Invariant Features: Why it works, what to do when it does not", talk at the *University of Padova*.

2007 "Viewpoint Invariant Features: Why it works, what to do when it does not", talk at te *University of San Diego (UCSD)*.

2007 "Viewpoint Invariant Features: Why it works, what to do when it does not", talk at the *University of California at Berkeley*.

Publications

[1] Minghao Chen, Iro Laina, and Andrea Vedaldi. "Training-Free Layout Control with Cross-Attention Guidance". In: *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*. 2024. URL: https://silent-chen.github.io/layout-guidance/.

[2] Minghao Chen, Junyu Xie, Iro Laina, and Andrea Vedaldi. "SHAP-EDITOR: Instruction-guided Latent 3D Editing in Seconds". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[3] Abdullah Hamdi, Luke Melas-Kyriazi, Guocheng Qian, Jinjie Mai, Ruoshi Liu, Carl Vondrick, Bernard Ghanem, and Andrea Vedaldi. "GES: Generalized Exponential Splatting for Efficient Radiance Field Rendering". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[4] Tomas Jakab, Ruining Li, Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. "Farm3D: Learning Articulated 3D Animals by Distilling 2D Diffusion". In: *Proceedings of the International Conference on 3D Vision (3DV)*. 2024. URL: https://farm3d.github.io.

[5] Zizhang Li, Dor Litvak, Ruining Li, Yunzhi Zhang, Tomas Jakab, Christian Rupprecht, Shangzhe Wu, Andrea Vedaldi, and Jiajun Wu. "Learning the 3D Fauna of the Web". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[6] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, Natalia Neverova, Andrea Vedaldi, Oran Gafni, and Filippos Kokkinos. "IM-3D: Iterative Multiview Diffusion and Reconstruction for High-Quality 3D Generation". In: *arXiv preprint* abs/2402.08682 (2024).

[7] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. "Splatter Image: Ultra-Fast Single-View 3D Reconstruction". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[8] Chuanxia Zheng and Andrea Vedaldi. "Free3D: Consistent Novel View Synthesis without 3D Representation". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[9] Yash Sanjay Bhalgat, Iro Laina, Joao F. Henriques, Andrea Vedaldi, and Andrew Zisserman. "Contrastive Lift: 3D Object Instance Segmentation by Slow-Fast Contrastive Fusion". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2023. URL: https://www.robots.ox.ac.uk/~vgg/research/contrastive-lift/.

[10] Changan Chen, Alexander Richard, Roman Shapovalov, Vamsi Ithapu, Natalia Neverova, Kristen Grauman, and Andrea Vedaldi. "Novel-view Acoustic Synthesis". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://replay-dataset.github.io.

[11] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. "DynamicStereo: Consistent Dynamic Depth from Stereo Videos". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://dynamic-stereo.github.io.

[12] Animesh Karnewar, Andrea Vedaldi, Niloy Mitra, and David Novotny. "HoloFusion: Towards Photo-realistic 3D Generative Modeling". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2023.

[13] Animesh Karnewar, Andrea Vedaldi, David Novotny, and Niloy Mitra. "HoloDiffusion: training a 3D diffusion model using 2D Images". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://holodiffusion.github.io.

[14] Yaoyao Liu, Bernt Schiele, Andrea Vedaldi, and Christian Rupprecht. "Continual Detection Transformer for Incremental Object Detection". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://lyy.mpi-inf.mpg.de/CL-DETR/.

[15] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. "RealFusion: 360 Reconstruction of Any Object from a Single Image". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://lukemelas.github.io/realfusion/.

[16] Luke Melas-Kyriazi, Christian Rupprecht, and Andrea Vedaldi. "PC2: Projection-Conditioned Point Cloud Diffusion for Single-Image 3D Reconstruction". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://lukemelas.github.io/projection-conditioned-point-cloud-diffusion/.

[17] Ignacio Rocco, Iurii Makarov, Filippos Kokkinos, David Novotny, Ben Graham, Natalia Neverova, and Andrea Vedaldi. "Replay: Multi-modal Multi-view Acted Videos for Casual Holography". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2023. URL: https://replay-dataset.github.io/.

[18] Ignacio Rocco, Iurii Makarov, Filippos Kokkinos, David Novotný, Benjamin Graham, Natalia Neverova, and Andrea Vedaldi. "Real-time volumetric rendering of dynamic humans". In: *arXiv preprint* abs/2303.11898 (2023). URL: https://real-time-humans.github.io.

[19] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. "Learning Universal Semantic Correspondences with No Supervision and Automatic Data Curation". In: *Proc. IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*. 2023.

[20] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. "What does CLIP know about a red circle? Visual prompt engineering for VLMs". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2023. URL: https://www.robots.ox.ac.uk/~vgg/publications/2023/Shtedritski23/.

[21] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, and Yaniv Taigman. "Text-To-4D Dynamic Scene Generation". In: *Proceedings of the International Conference on Machine Learning (ICML)*. 2023. URL: https://make-a-video3d.github.io.

[22] Samarth Sinha, Roman Shapovalov, Jeremy Reizenstein, Ignacio Rocco, Natalia Neverova, Andrea Vedaldi, and David Novotný. "Common Pets in 3D: Dynamic New-View Synthesis of Real-Life Deformable Categories". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://cop3d.github.io.

[23] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. "Viewset Diffusion: (0-)Image-Conditioned 3D Generative Models from 2D Data". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2023. URL: https://szymanowiczs.github.io/viewset-diffusion.

[24] Vadim Tschernezki, Ahmad Darkhalil, Zhifan Zhu, David Fouhey, Iro Laina, Diane Larlus, Dima Damen, and Andrea Vedaldi. "EPIC Fields: Marrying 3D Geometry and Video Understanding". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS): Track on Datasets and Benchmarks*. 2023. URL: http://epic-kitchens.github.io/epic-fields.

[25] Sagar Vaze, Andrea Vedaldi, and Andrew Zisserman. "No Representation Rules Them All in Category Discovery". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2023.

[26] Felix Wimbauer, Nan Yang, Christian Rupprecht, and Daniel Cremers. "Behind the Scenes: Density Fields for Single View Reconstruction". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://fwmb.github.io/bts/.

[27] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. "MagicPony: Learning Articulated 3D Animals in the Wild". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. URL: https://3dmagicpony.github.io.

[28] Chuanxia Zheng and Andrea Vedaldi. "Online Clustered Codebook". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2023.

[29] Triantafyllos Afouras, Yuki Asano, Francois Fagan, Andrea Vedaldi, and Florian Metze. "Self-supervised object detection from audio-visual correspondence". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022.

[30] Andrew Brown, Cheng-Yang Fu, Omkar M. Parkhi, Tamara L. Berg, and Andrea Vedaldi. "End-to-End Visual Editing with a Generatively Pre-Trained Artist". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2022. URL: https://www.robots.ox.ac.uk/~abrown/E2EVE/.

[31] Eldar Insafutdinov, Dylan Campbell, João F. Henriques, and Andrea Vedaldi. "SNeS: Learning Probably Symmetric Neural Surfaces from Incomplete Data". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2022. URL: https://www.robots.ox.ac.uk/~vgg/research/snes/.

[32] Xu Ji, Razvan Pascanu, R. Devon Hjelm, Balaji Lakshminarayanan, and Andrea Vedaldi. "Test Sample Accuracy Scales with Training Sample Density in Neural Networks". In: *Proceedings of Conference on Lifelong Learning Agents (CoLLA)*. 2022.

[33] Laurynas Karazija, Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. "Unsupervised Multi-object Segmentation by Predicting Probable Motion Patterns". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2022. URL: https://www.robots.ox.ac.uk/~vgg/research/ppmp/.

[34] Laurynas Karazija, Subhabrata Choudhury, Iro Laina, Andrea Vedaldi, and Christian Rupprecht. "Guess What Moves: Unsupervised Video and Image Segmentation by Anticipating Motion". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2022. URL: https://www.robots.ox.ac.uk/~vgg/research/gwm/.

[35] Iro Laina, Yuki M. Asano, and Andrea Vedaldi. "Measuring the Interpretability of Unsupervised Representations via Quantized Reversed Probing". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2022. URL: https://www.robots.ox.ac.uk/~vgg/publications/2022/Laina22/.

[36] Robert McCraith, Eldar Insafutdinov, Lukas Neumann, and Andrea Vedaldi. "Direct LiDAR-based object detector training from automated 2D detections". In: *Proceedings of NeurIPS Machine Learning for the Developing World Workshop*. 2022.

[37] Robert McCraith, Eldar Insafutdinov, Lukás Neumann, and Andrea Vedaldi. "Lifting 2D Object Locations to 3D by Discounting LiDAR Outliers across Objects and Views". In: *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*. 2022.

[38] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. "Deep Spectral Methods: A Surprisingly Strong Baseline for Unsupervised Semantic Segmentation and Localization". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022. URL: https://lukemelas.github.io/deep-spectral-segmentation/.

[39] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. "Finding an Unsupervised Image Segmenter in each of your Deep Generative Models". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2022. URL: https://github.com/lukemelas/unsupervised-image-segmentation.

[40] David Novotny, Ignacio Rocco, Samrath Sinha, Alexandre Carlier, Gael Kerchenbaum, Roman Shapovalov, Nikita Smetanin, Natalia Neverova, and Andrea Vedaldi. "Keypoint Transporter for Reconstructing Non-rigid 3D Shapes in Videos". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022.

[41] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. "Neural Feature Fusion Fields: 3D Distillation of Self-Supervised 2D Image Representation". In: *Proceedings of the International Conference on 3D Vision (3DV)*. 2022. URL: https://www.robots.ox.ac.uk/~vadim/n3f/.

[42] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. "Generalized Category Discovery". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022. URL: https://www.robots.ox.ac.uk/~vgg/research/gcd/.

[43] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. "Open-Set Recognition: A Good Closed-Set Classifier is All You Need". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2022. URL: https://github.com/sgvaze/osr_closed_set_all_you_need.

[44] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. "BANMo: Building Animatable 3D Neural Models from Many Casual Videos". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022. URL: https://banmo-www.github.io.

[45] Yuki M. Asano, Christian Rupprecht, Andrew Zisserman, and Andrea Vedaldi. "PASS: An ImageNet replacement for self-supervised pretraining without human". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS): Track on Datasets and Benchmarks* (2021). URL: https://www.robots.ox.ac.uk/~vgg/data/pass/.

[46] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. "Audio-Visual Synchronisation in the Wild". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2021.

[47] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. "The Curious Layperson: Fine-Grained Image Recognition without Expert Labels". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2021. URL: https://www.robots.ox.ac.uk/~vgg/research/clever/.

[48] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. "Unsupervised Part Discovery from Contrastive Reconstruction". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2021. URL: https://www.robots.ox.ac.uk/~vgg/research/unsup-parts/.

[49] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. "NeuroMorph: Unsupervised Shape Interpolation and Correspondence in One Go". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021.

[50] Weikang Gong, Christian F. Beckmann, Andrea Vedaldi, Stephen M. Smith, and Han Peng. "Optimising a Simple Fully Convolutional Network (SFCN) for accurate brain age prediction in the PAC 2019 challenge". In: *Front. Psychiatry* 12 (2021).

[51] Oliver Groth, Chia-Man Hung, Andrea Vedaldi, and Ingmar Posner. "Goal-Conditioned End-to-End Visuomotor Control for Versatile Skill Primitives". In: *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*. 2021. URL: https://github.com/ogroth/geeco/.

[52] Kai Han, Sylvestre-Alvise Rebuffi, Sébastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. "AutoNovel: Automatically Discovering and Learning Novel Visual Categories". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021). URL: https://www.robots.ox.ac.uk/~vgg/research/auto_novel/.

[53] Philipp Henzler, Jeremy Reizenstein, Patrick Labatut, Roman Shapovalov, Tobias Ritschel, Andrea Vedaldi, and David Novotny. "Unsupervised Learning of 3D Object Categories from Videos in the Wild". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021.

[54] Robert McCraith, Lukas Neumann, and Andrea Vedaldi. "Real Time Monocular Vehicle Velocity Estimation Using Synthetic Data". In: *Proc. IEEE Intelligent Vehicles Symposium*. 2021.

[55] Natalia Neverova, Artsiom Sanakoyeu, Patrick Labatut, David Novotny, and Andrea Vedaldi. "Discovering Relationships between Object Categories via Universal Canonical Maps". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021.

[56] Mandela Patrick, Yuki M. Asano, Bernie Huang, Ishan Misra, Florian Metze, João F. Henriques, and Andrea Vedaldi. "Space-time crop and Attend: Improving cross-modal video representation learning". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2021.

[57] Mandela Patrick, Yuki M. Asano, Polina Kuznetsova, Ruth Fong, João F. Henriques, Geoffrey Zweig, and Andrea Vedaldi. "On Compositions of Transformations in Contrastive Self-Supervised Learning". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2021.

[58] Mandela Patrick, Dylan Campbell, Yuki Markus Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and João F. Henriques. "Keeping Your Eye on the Ball: Trajectory Attention in Video Transformers". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2021.

[59] Mandela Patrick, Po-Yao Huang, Yuki Markus Asano, Florian Metze, Alexander G. Hauptmann, João F. Henriques, and Andrea Vedaldi. "Support-set bottlenecks for video-text representation learning". In: *iclr*. 2021.

[60] Roman Shapovalov, David Novotny, Benjamin Graham, Patrick Labatut, and Andrea Vedaldi. "DensePose 3D: Lifting Canonical Surface Maps of Articulated Objects to the Third Dimension". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2021.

[61] Vadim Tschernezki, Diane Larlus, and Andrea Vedaldi. "NeuralDiff: Segmenting 3D objects that move in egocentric videos". In: *Proceedings of the International Conference on 3D Vision (3DV)*. 2021. URL: https://www.robots.ox.ac.uk/~vadim/neuraldiff/.

[62] Dan Xu, Andrea Vedaldi, and João F. Henriques. "Moving SLAM: Fully Unsupervised Deep Learning in Non-Rigid Scenes". In: *Proc. IEEE Intl. Conf. on Intelligent Robots an Systems (IROS)*. 2021.

[63] Yuki M. Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. "Labelling unlabelled videos from scratch with multi-modal self-supervision". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[64] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. "A critical analysis of self-supervision, or what we can learn from a single image". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2020.

[65] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. "Self-labelling via simultaneous clustering and representation learning". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2020.

[66] Benjamin Biggs, Andrea Vedaldi, and David Novotný. "3D Multi-bodies: Fitting Sets of Plausible 3D Human Models to Ambiguous Image Data". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[67] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. "VGGSound: A Large-scale Audio-Visual Dataset". In: *International Conference on Acoustics, Speech, and Signal Processing (ICASS)*. 2020.

[68] Sébastien Ehrhardt, Oliver Groth, Aron Monszpart, Martin Engelcke, Ingmar Posner, Niloy J. Mitra, and Andrea Vedaldi. "RELATE: Physically Plausible Multi-Object Scene Synthesis Using Structured Latent Spaces". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[69] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. "Automatically Discovering and Learning New Visual Categories with Ranking Statistics". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2020.

[70] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. "Self-supervised Learning of Interpretable Keypoints from Unlabelled Videos". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.

[71] Iro Laina, Ruth C. Fong, and Andrea Vedaldi. "Quantifying Learnability and Describability of Visual Concepts Emerging in Representation Learning". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[72] Natalia Neverova, David Novotný, and Andrea Vedaldi. "Continuous Surface Embeddings". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[73] David Novotný, Roman Shapovalov, and Andrea Vedaldi. "Canonical 3D Deformer Maps: Unifying parametric and non-parametric methods for dense weakly-supervised category reconstruction". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2020.

[74] Sylvestre-Alvise Rebuffi, Ruth C. Fong, Xu Ji, and Andrea Vedaldi. "There and Back Again: Revisiting Backpropagation Saliency Methods". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.

[75] Artsiom Sanakoyeu, Vasil Khalidov, Maureen S. McCarthy, Andrea Vedaldi, and Natalia Neverova. "Transferring Dense Pose to Proximal Animal Classes". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.

[76] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. "Deep Image Prior". In: *International Journal of Computer Vision (IJCV)* 128.7 (2020), pp. 1867–1888.

[77] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. "Unsupervised Learning of Probably Symmetric Deformable 3D Objects from Images in the Wild". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020. URL: https://elliottwu.com/projects/unsup3d/.

[78] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, Tinne Tuytelaars, Jiri Matas, and Krystian Mikolajczyk. "HPatches: A benchmark and evaluation of handcrafted and learned local descriptors". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).

[79] Luca Bertinetto, João F. Henriques, Philip H. S. Torr, and Andrea Vedaldi. "Meta-learning with differentiable closed-form solvers". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2019. URL: https://www.robots.ox.ac.uk/~luca/r2d2.html.

[80] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. "AutoCorrect: Deep Inductive Alignment of Noisy Geometric Annotations". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2019.

[81] Sébastien Ehrhardt, Aron Monszpart, Niloy J. Mitra, and Andrea Vedaldi. "Taking visual motion prediction to new heightfields". In: *Computer Vision and Image Understanding* (2019).

[82] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. "Understanding Deep Networks via Extremal Perturbations and Smooth Masks". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[83] Fabian B. Fuchs, Oliver Groth, Adam R. Kosiorek, Alex Bewley, Markus Wulfmeier, Andrea Vedaldi, and Ingmar Posner. "Scrutinizing and De-Biasing Intuitive Physics with Neural Stethoscopes". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2019.

[84] Kai Han, Andrea Vedaldi, and Andrew Zisserman. "Learning to Discover Novel Visual Categories via Deep Transfer Clustering". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[85] João F. Henriques, Sébastien Ehrhardt, Samuel Albanie, and Andrea Vedaldi. "Small step and giant leaps: Minimal Netwon solvers for Deep Learning". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[86] Xu Ji, João F. Henriques, and Andrea Vedaldi. "Invariant Information Clustering for Unsupervised Image Classification and Segmentation". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[87] Karel Lenc and Andrea Vedaldi. "Understanding Image Representations by Measuring Their Equivariance and Equivalence". In: *International Journal of Computer Vision (IJCV)* 127.5 (2019).

[88] Lukás Neumann, Andrew Zisserman, and Andrea Vedaldi. "Future Event Prediction: If and When". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2019.

[89] Natalia Neverova, David Novotný, and Andrea Vedaldi. "Correlated Uncertainty for Learning Dense Correspondences from Noisy Labels". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2019.

[90] Natalia Neverova, James Thewlis, Riza Alp Güler, Iasonas Kokkinos, and Andrea Vedaldi. "Slim DensePose: Thrifty Learning From Sparse Annotations and Motion Cues". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019.

[91] David Novotný, Nikhila Ravi, Benjamin Graham, Natalia Neverova, and Andrea Vedaldi. "C3DPO: Canonical 3D Pose Networks for Non-Rigid Structure From Motion". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[92] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller, eds. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 2019.

[93] James Thewlis, Samuel Albanie, Hakan Bilen, and Andrea Vedaldi. "Unsupervised learning of landmarks by descriptor vector exchange". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2019.

[94] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. "Fixing the train-test resolution discrepancy". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2019.

[95] Samuel Albanie, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. "Emotion Recognition in Speech using Cross-Modal Transfer in the Wild". In: *Proceedings of the 26th annual ACM international conference on Multimedia*. 2018.

[96] Hakan Bilen, Basura Fernando, Efstratios Gavves, and Andrea Vedaldi. "Action Recognition with Dynamic Image Networks". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40.12 (2018).

[97] Sébastien Ehrhardt, Aron Monszpart, Niloy J. Mitra, and Andrea Vedaldi. "Unsupervised Intuitive Physics from Visual Observations". In: *Proceedings of the Asian Conference on Computer Vision (ACCV)*. 2018.

[98] Ruth Fong and Andrea Vedaldi. "Net2Vec: Quantifying and Explaining How Concepts Are Encoded by Filters in Deep Neural Networks". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.

[99] Oliver Groth, Fabian B. Fuchs, Ingmar Posner, and Andrea Vedaldi. "ShapeStacks: Learning Vision-Based Physical Intuition for Generalised Object Stacking". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018. URL: https://shapestacks.robots.ox.ac.uk.

[100] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. "Inductive Visual Localisation: Factorised Training for Superior Generalisation". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2018.

[101] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. "Learning to read by spelling: Towards unsupervised text recognition". In: *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*. 2018.

[102] João F. Henriques and Andrea Vedaldi. "MapNet: An Allocentric Spatial Memory for Mapping Environments". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.

[103] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Andrea Vedaldi. "Gather-Excite: Exploiting Feature Context in Convolutional Neural Networks". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2018.

[104] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. "Unsupervised Learning of Object Landmarks through Conditional Image Generation". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2018.

[105] Maria Klodt and Andrea Vedaldi. "Supervising the New with the Old: Learning SFM from SFM". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.

[106] Karel Lenc and Andrea Vedaldi. "Large scale evaluation of local image feature detectors on homography datasets". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2018.

[107] Aravindh Mahendran, James Thewlis, and Andrea Vedaldi. "Cross Pixel Optical-Flow Similarity for Self-supervised Learning". In: *Proceedings of the Asian Conference on Computer Vision (ACCV)*. 2018.

[108] Lukás Neumann, Michelle Karg, Shanshan Zhang, Christian Scharfenberger, Eric Piegert, Sarah Mistr, Olga Prokofyeva, Robert Thiel, Andrea Vedaldi, Andrew Zisserman, and Bernt Schiele. "NightOwls: A Pedestrians at Night Dataset". In: *Proceedings of the Asian Conference on Computer Vision (ACCV)*. 2018.

[109] Lukás Neumann and Andrea Vedaldi. "Tiny People Pose". In: *Proceedings of the Asian Conference on Computer Vision (ACCV)*. 2018.

[110] Lukás Neumann, Andrew Zisserman, and Andrea Vedaldi. "Relaxed Softmax: Efficient Confidence Auto-Calibration for Safe Pedestrian Detection". In: *NeurIPS Workshop on Machine Learning for Intelligent Transportation Systems*. 2018.

[111] David Novotný, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. "Self-Supervised Learning of Geometrically Stable Features Through Probabilistic Introspection". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.

[112] David Novotný, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. "Semi-convolutional Operators for Instance Segmentation". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.

[113] David Novotný, Diane Larlus, and Andrea Vedaldi. "Capturing the Geometry of Object Categories from Video Supervision". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018).

[114] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. "Efficient Parametrization of Multi-Domain Deep Neural Networks". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018. URL: https://github.com/srebuffi/residual_adapters.

[115] James Thewlis, Hakan Bilen, and Andrea Vedaldi. "Modelling and unsupervised learning of symmetric deformable object categories". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2018.

[116] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. "Deep Image Prior". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018. URL: https://dmitryulyanov.github.io/deep_image_prior.

[117] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. "It Takes (Only) Two: Adversarial Generator-Encoder Networks". In: *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*. 2018.

[118] Jack Valmadre, Luca Bertinetto, João F. Henriques, Ran Tao, Andrea Vedaldi, Arnold W. M. Smeulders, Philip H. S. Torr, and Efstratios Gavves. "Long-Term Tracking in the Wild: A Benchmark". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.

[119] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. "HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

[120] Ruth C. Fong and Andrea Vedaldi. "Interpretable Explanations of Black Boxes by Meaningful Perturbation". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2017.

[121] João F. Henriques and Andrea Vedaldi. "Warped Convolutions: Efficient Invariance to Spatial Transformations". In: *Proceedings of the International Conference on Machine Learning (ICML)*. 2017.

[122] David Novotný, Diane Larlus, and Andrea Vedaldi. "AnchorNet: A Weakly Supervised Network to Learn Geometry-Sensitive Features for Semantic Matching". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

[123] David Novotný, Diane Larlus, and Andrea Vedaldi. "Generalizing Semantic Part Detectors Across Domains". In: *Domain Adaptation in Computer Vision Applications*. Springer, Cham, 2017, pp. 259–273.

[124] David Novotný, Diane Larlus, and Andrea Vedaldi. "Learning 3D Object Categories by Looking Around Them". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2017.

[125] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. "Learning multiple visual domains with residual adapters". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2017. URL: https://github.com/srebuffi/residual_adapters.

[126] James Thewlis, Hakan Bilen, and Andrea Vedaldi. "Unsupervised learning of object frames by dense equivariant image labelling". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2017.

[127] James Thewlis, Hakan Bilen, and Andrea Vedaldi. "Unsupervised Learning of Object Landmarks by Factorized Spatial Embeddings". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2017.

[128] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. "Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

[129] Jack Valmadre, Luca Bertinetto, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. "End-to-End Representation Learning for Correlation Filter Based Tracking". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

[130] Samuel Albanie and Andrea Vedaldi. "Learning Grimaces by Watching TV". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2016. URL: http://www.robots.ox.ac.uk/~vgg/data/facevalue.

[131] Luca Bertinetto, João F. Henriques, Jack Valmadre, Philip H. S. Torr, and Andrea Vedaldi. "Learning feed-forward one-shot learners". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2016.

[132] Luca Bertinetto, Jack Valmadre, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. "Fully-Convolutional Siamese Networks for Object Tracking". In: *ECCV Workshop on Computer Vision on Visual Object Tracking*. 2016.

[133] Hakan Bilen, Basura Fernando, Efstratios Gavves, Andrea Vedaldi, and Stephen Gould. "Dynamic Image Networks for Action Recognition". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.

[134] Hakan Bilen and Andrea Vedaldi. "Integrated perception with recurrent multi-task neural networks". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2016.

[135] Hakan Bilen and Andrea Vedaldi. "Weakly Supervised Deep Detection Networks". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.

[136] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, and Andrea Vedaldi. "Deep Filter Banks for Texture Recognition, Description, and Segmentation". In: *International Journal of Computer Vision (IJCV)* 118.1 (2016).

[137] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. "Synthetic Data for Text Localisation in Natural Images". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.

[138] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Reading Text in the Wild with Convolutional Neural Networks". In: *International Journal of Computer Vision (IJCV)* 116.1 (2016).

[139] Karel Lenc and Andrea Vedaldi. "Learning Covariant Feature Detectors". In: *Proceedings of the ECCV workshop on Geometry Meets Deep Learning*. 2016. URL: https://github.com/lenck/ddet.

[140] Aravindh Mahendran and Andrea Vedaldi. "Salient Deconvolutional Networks". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2016.

[141] Aravindh Mahendran and Andrea Vedaldi. "Visualizing Deep Convolutional Neural Networks Using Natural Pre-images". In: *International Journal of Computer Vision (IJCV)* 120.3 (2016).

[142] David Novotný, Diane Larlus, and Andrea Vedaldi. "I Have Seen Enough: Transferring Parts Across Categories". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2016.

[143] David Novotný, Diane Larlus, and Andrea Vedaldi. "Learning the semantic structure of objects from Web supervision". In: *Proceedings of the ECCV workshop on Geometry Meets Deep Learning*. 2016.

[144] James Thewlis, Shuai Zheng, Philip H. S. Torr, and Andrea Vedaldi. "Fully-trainable deep matching". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2016.

[145] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S. Lempitsky. "Texture Networks: Feed-forward Synthesis of Textures and Stylized Images". In: *Proceedings of the International Conference on Machine Learning (ICML)*. 2016. URL: http://likemo.net/.

[146] Mircea Cimpoi, Subhransu Maji, and Andrea Vedaldi. "Deep filter banks for texture recognition and segmentation". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015. URL: http://www.robots.ox.ac.uk/~vgg/data/dtd/.

[147] Josep M. Gonfaus, Marco Pedersoli, Jordi Gonzàlez, Andrea Vedaldi, and F. Xavier Roca. "Factorized appearances for object detection". In: *Computer Vision and Image Understanding* 138 (2015), pp. 92–101.

[148] Karel Lenc and Andrea Vedaldi. "R-CNN minus R". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2015.

[149] Karel Lenc and Andrea Vedaldi. "Understanding image representations by measuring their equivariance and equivalence". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.

[150] Aravindh Mahendran and Andrea Vedaldi. "Understanding deep image representations by inverting them". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.

[151] Sobhan Naderi Parizi, Andrea Vedaldi, Andrew Zisserman, and Pedro F. Felzenszwalb. "Automatic Discovery and Optimization of Parts for Image Classification". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2015.

[152] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. "Deep Face Recognition". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2015.

[153] Andrea Vedaldi and Karel Lenc. "MatConvNet: Convolutional Neural Networks for MATLAB". In: *Proceedings of the 25th annual ACM international conference on Multimedia*. 2015.

[154] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Return of the Devil in the Details: Delving Deep into Convolutional Nets". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2014. URL: http://www.robots.ox.ac.uk/~vgg/research/deep_eval/.

[155] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. "Describing Textures in the Wild". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014. URL: http://www.robots.ox.ac.uk/~vgg/data/dtd/.

[156] M. Jaderberg, K. Simonyan, Andrea Vedaldi, and A. Zisserman. "Synthetic Data and Artificial Neural Networks for Natural Scene Text Recognition". In: *NeurIPS Deep Learning Workshop*. 2014.

[157] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. "Deep Features for Text Spotting". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2014.

[158] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. "Speeding up Convolutional Neural Networks with Low Rank Expansions". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2014.

[159] Omkar M. Parkhi, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "A Compact and Discriminative Face Track Descriptor". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.

[160] Marco Pedersoli, Andrea Vedaldi, Jordi Gonzàlez, and F. Xavier Roca. "A coarse-to-fine approach for fast deformable object detection". In: *Pattern Recognition* (2014).

[161] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2014.

[162] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Learning Local Feature Descriptors Using Convex Optimisation". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36.8 (2014).

[163] Andrea Vedaldi, Siddharth Mahendran, Stavros Tsogkas, Subhransu Maji, Ross B. Girshick, Juho Kannala, Esa Rahtu, Iasonas Kokkinos, Matthew B. Blaschko, David J. Weiss, Ben Taskar, Karen Simonyan, Naomi Saphra, and Sammy Mohamed. "Understanding Objects in Detail with Fine-Grained Attributes". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014. URL: http://www.robots.ox.ac.uk/~vgg/data/oid/.

[164] Mayank Juneja, Andrea Vedaldi, C. V. Jawahar, and Andrew Zisserman. "Blocks That Shout: Distinctive Parts for Scene Classification". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2013.

[165] Karen Simonyan, Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. "Fisher Vector Faces in the Wild". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2013.

[166] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Deep Fisher Networks for Large-Scale Image Classification". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2013. URL: http://www.robots.ox.ac.uk/~vgg/software/face_desc/.

[167] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. "On-the-fly specific person retrieval". In: *Proceedings of the Int. Workshop on Image Analysis for Multimedia Interactive Services*. 2012.

[168] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. "Cats and dogs". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.

[169] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. "Descriptor Learning Using Convex Optimisation". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2012.

[170] Andrea Vedaldi and Andrew Zisserman. "Efficient Additive Kernels via Explicit Feature Maps". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34.3 (2012).

[171] Andrea Vedaldi and Andrew Zisserman. "Self-similar Sketch". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2012.

[172] Andrea Vedaldi and Andrew Zisserman. "Sparse kernel approximations for efficient classification and detection". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.

[173] Ken Chatfield, Victor S. Lempitsky, Andrea Vedaldi, and Andrew Zisserman. "The devil is in the details: an evaluation of recent feature encoding methods". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2011. URL: http://www.robots.ox.ac.uk/~vgg/research/encoding_eval/.

[174] Victor S. Lempitsky, Andrea Vedaldi, and Andrew Zisserman. "Pylon Model for Semantic Segmentation". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2011.

[175] Omkar M. Parkhi, Andrea Vedaldi, C. V. Jawahar, and Andrew Zisserman. "The truth about cats and dogs". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2011.

[176] Marco Pedersoli, Andrea Vedaldi, and Jordi Gonzàlez. "A coarse-to-fine approach for fast deformable object detection". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2011.

[177] Andrea Vedaldi, Matthew B. Blaschko, and Andrew Zisserman. "Learning equivariant structured output SVM regressors". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2011.

[178] Matthew B. Blaschko, Andrea Vedaldi, and Andrew Zisserman. "Simultaneous Object Detection and Ranking with Weak Supervision". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2010.

[179] Andrea Vedaldi and Brian Fulkerson. "VLFeat: an open and portable library of computer vision algorithms". In: *Proceedings of the 18th annual ACM international conference on Multimedia*. 2010.

[180] Andrea Vedaldi, Haibin Ling, and Stefano Soatto. "Knowing a Good Feature When You See It: Ground Truth and Methodology to Evaluate Local Features for Recognition". In: *Computer Vision: Detection, Recognition and Reconstruction*. Vol. 285. Springer, 2010, pp. 27–49.

[181] Andrea Vedaldi and Andrew Zisserman. "Efficient additive kernels via explicit feature maps". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2010.

[182] Sreekanth Vempati, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. "Generalized RBF feature maps for Efficient Detection". In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2010.

[183] Brian Fulkerson, Andrea Vedaldi, and Stefano Soatto. "Class segmentation and object localization with superpixel neighborhoods". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2009.

[184] Andrea Vedaldi, Varun Gulshan, Manik Varma, and Andrew Zisserman. "Multiple kernels for object detection". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2009.

[185] Andrea Vedaldi and A. Zisserman. "Structured output regression for detection with partial occulsion". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2009.

[186] Brian Fulkerson, Andrea Vedaldi, and Stefano Soatto. "Localizing Objects with Smart Dictionaries". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2008.

[187] Andrea Vedaldi. "Invariant Representations and Learning for Computer Vision". PhD thesis. University of California at Los Angeles, 2008.

[188] Andrea Vedaldi, Gregorio Guidi, and Stefano Soatto. "Joint data alignment up to (lossy) transformations". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2008.

[189] Andrea Vedaldi and Stefano Soatto. "Quick Shift and Kernel Methods for Mode Seeking". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2008.

[190] Andrea Vedaldi and Stefano Soatto. "Relaxed Matching Kernels for Object Recognition". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2008.

[191] Eagle Jones, Andrea Vedaldi, and Stefano Soatto. "Inertial structure from motion with autocalibration". In: *Proceedings of ICCV Workshop on Dynamical Vision*. 2007.

[192] Andrew Rabinovich, Andrea Vedaldi, Carolina Galleguillos, Eric Wiewiora, and Serge J. Belongie. "Objects in Context". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2007.

[193] Andrea Vedaldi. *An open implementation of the SIFT detector and descriptor*. Tech. rep. 070012. UCLA CSD, 2007.

[194] Andrea Vedaldi, Paolo Favaro, and Enrico Grisan. "Boosting Invariance and Efficiency in Supervised Learning". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2007.

[195]   Andrea Vedaldi, Gregorio Guidi, and Stefano Soatto. "Moving Forward in Structure From Motion".
         In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
         2007.

[196]   Andrea Vedaldi and Stefano Soatto. "A Complexity-Distrotion Approach to Joint Pattern
         Alignment". In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*.
         2007.

[197]   Andrea Vedaldi and Stefano Soatto. "Local Features, All Grown Up". In: *Proceedings of the IEEE
         Conference on Computer Vision and Pattern Recognition (CVPR)*. 2006.

[198]   Andrea Vedaldi and Stefano Soatto. "Viewpoint Induced Deformation Statistics and the Design of
         Viewpoint Invariant Features: Singularities and Occlusions". In: *Proceedings of the European
         Conference on Computer Vision (ECCV)*. 2006, p. 374.

[199]   Andrea Vedaldi, Hailin Jin, Paolo Favaro, and Stefano Soatto. "KALMANSAC: Robust Filtering by
         Consensus". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. 2005.

[200]   Andrea Vedaldi and Stefano Soatto. "Features for Recognition: Viewpoint Invariance for
         Non-Planar Scenes". In: *Proceedings of the International Conference on Computer Vision (ICCV)*.
         2005.

[201]   Andrea Vedaldi and Stefano Soatto. *On Viewpoint Invariance for Non-Planar Scenes*. Tech. rep.
         050012. UCLA CSD, 2005.