

Andrew Zisserman

List of Publications by Citations

Source: <https://exaly.com/author-pdf/6676312/andrew-zisserman-publications-by-citations.pdf>
Version: 2024-04-03

This document has been generated based on the publications and citations recorded by exaly.com. For the latest version of this publication list, visit the link given above.
The third column is the impact factor (IF) of the journal, and the fourth column is the number of citations of the article.

107 papers	25,462 citations	49 h-index	117 g-index
117 ext. papers	34,387 ext. citations	6.9 avg, IF	7.78 L-index

#	Paper	IF	Citations
107	The Pascal Visual Object Classes (VOC) Challenge. <i>International Journal of Computer Vision</i> , 2010 , 88, 303-338	10.6	7456
106	Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset 2017 ,		1965
105	Object retrieval with large vocabularies and fast spatial matching 2007 ,		1598
104	Deep Face Recognition 2015 ,		1422
103	Convolutional Two-Stream Network Fusion for Video Action Recognition 2016 ,		1092
102	Return of the Devil in the Details: Delving Deep into Convolutional Nets 2014 ,		1048
101	Automated Flower Classification over a Large Number of Classes 2008 ,		814
100	A Statistical Approach to Texture Classification from Single Images. <i>International Journal of Computer Vision</i> , 2005 , 62, 61-81	10.6	810
99	Lost in quantization: Improving particular object retrieval in large scale image databases 2008 ,		804
98	2018 ,		706
97	Reading Text in the Wild with Convolutional Neural Networks. <i>International Journal of Computer Vision</i> , 2016 , 116, 1-20	10.6	581
96	Synthetic Data for Text Localisation in Natural Images 2016 ,		520
95	The devil is in the details: an evaluation of recent feature encoding methods 2011 ,		490
94	Total Recall: Automatic Query Expansion with a Generative Feature Model for Object Retrieval 2007 ,		487
93	All About VLAD 2013 ,		354
92	VoxCeleb2: Deep Speaker Recognition		313
91	Non-uniform Deblurring for Shaken Images. <i>International Journal of Computer Vision</i> , 2012 , 98, 168-186	10.6	286

90	VoxCeleb: A Large-Scale Speaker Identification Dataset		272
89	A framework for spatiotemporal control in the tracking of visual contours. <i>International Journal of Computer Vision</i> , 1993 , 11, 127-145	10.6	235
88	Flowing ConvNets for Human Pose Estimation in Videos 2015 ,		234
87	Cats and dogs 2012 ,		222
86	Geodesic star convexity for interactive image segmentation 2010 ,		191
85	Look, Listen and Learn 2017 ,		187
84	Tabula rasa: Model transfer for object category detection 2011 ,		182
83	An Exemplar Model for Learning Object Classes 2007 ,		151
82	Microscopy cell counting and detection with fully convolutional regression networks. <i>Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization</i> , 2018 , 6, 283-292	0.9	150
81	Symbiotic Segmentation and Part Localization for Fine-Grained Categorization 2013 ,		140
80	Creating Architectural Models from Images. <i>Computer Graphics Forum</i> , 1999 , 18, 39-50	2.4	134
79	Non-uniform deblurring for shaken images 2010 ,		131
78	Triangulation Embedding and Democratic Aggregation for Image Search 2014 ,		127
77	Who are you? Learning person specific classifiers from video 2009 ,		89
76	Emotion Recognition in Speech using Cross-Modal Transfer in the Wild 2018 ,		83
75	Diagnostically relevant facial gestalt information from ordinary photos. <i>ELife</i> , 2014 , 3, e02020	8.9	81
74	A Compact and Discriminative Face Track Descriptor 2014 ,		79
73	Deblurring Shaken and Partially Saturated Images. <i>International Journal of Computer Vision</i> , 2014 , 110, 185-201	10.6	79

72	Counting in the Wild. <i>Lecture Notes in Computer Science</i> , 2016 , 483-498	0.9	75
71	End-to-End Learning of Visual Representations From Uncurated Instructional Videos 2020 ,		74
70	Objects that Sound. <i>Lecture Notes in Computer Science</i> , 2018 , 451-466	0.9	73
69	ISSLS PRIZE IN BIOENGINEERING SCIENCE 2017: Automation of reading of radiological features from magnetic resonance images (MRIs) of the lumbar spine without human intervention is comparable with an expert radiologist. <i>European Spine Journal</i> , 2017 , 26, 1374-1383	2.7	72
68	2019 ,		67
67	Utterance-level Aggregation for Speaker Recognition in the Wild 2019 ,		66
66	Learning an Alphabet of Shape and Appearance for Multi-Class Object Detection. <i>International Journal of Computer Vision</i> , 2008 , 80, 16-44	10.6	64
65	Object Level Grouping for Video Shots. <i>International Journal of Computer Vision</i> , 2006 , 67, 189-210	10.6	61
64	Seeing Voices and Hearing Faces: Cross-Modal Biometric Matching 2018 ,		59
63	SpineNet: Automated classification and evidence visualization in spinal MRIs. <i>Medical Image Analysis</i> , 2017 , 41, 63-73	15.4	58
62	Chimpanzee face recognition from videos in the wild using deep learning. <i>Science Advances</i> , 2019 , 5, eaaw0736	14.3	57
61	Crystal nucleation in metallic alloys using x-ray radiography and machine learning. <i>Science Advances</i> , 2018 , 4, eaar4004	14.3	55
60	Object Mining Using a Matching Graph on Very Large Image Collections 2008 ,		52
59	Interactive Object Counting. <i>Lecture Notes in Computer Science</i> , 2014 , 504-518	0.9	51
58	Sparse kernel approximations for efficient classification and detection 2012 ,		49
57	Geometric Latent Dirichlet Allocation on a Matching Graph for Large-scale Image Datasets. <i>International Journal of Computer Vision</i> , 2011 , 95, 138-153	10.6	47
56	EPIC-Fusion: Audio-Visual Temporal Binding for Egocentric Action Recognition 2019 ,		47
55	Template Adaptation for Face Verification and Identification 2017 ,		46

54	Detecting overlapping instances in microscopy images using extremal region trees. <i>Medical Image Analysis</i> , 2016 , 27, 3-16	15.4	41
53	Fully-automated alignment of 3D fetal brain ultrasound to a canonical reference space using multi-task learning. <i>Medical Image Analysis</i> , 2018 , 46, 1-14	15.4	39
52	Deblurring shaken and partially saturated images 2011 ,		39
51	Upper Body Detection and Tracking in Extended Signing Sequences. <i>International Journal of Computer Vision</i> , 2011 , 95, 180-197	10.6	38
50	Non-contact physiological monitoring of preterm infants in the Neonatal Intensive Care Unit. <i>Npj Digital Medicine</i> , 2019 , 2, 128	15.7	37
49	Multi-Task Convolutional Neural Network for Patient Detection and Skin Segmentation in Continuous Non-Contact Vital Sign Monitoring 2017 ,		35
48	Out of Time: Automated Lip Sync in the Wild. <i>Lecture Notes in Computer Science</i> , 2017 , 251-263	0.9	35
47	Automatic and Efficient Human Pose Estimation for Sign Language Videos. <i>International Journal of Computer Vision</i> , 2014 , 110, 70-90	10.6	33
46	Resolving ambiguities in auto-calibration. <i>Philosophical Transactions Series A, Mathematical, Physical, and Engineering Sciences</i> , 1998 , 356, 1193-1211	3	33
45	Smooth object retrieval using a bag of boundaries 2011 ,		32
44	A Statistical Approach to Texture Classification from Single Images. <i>International Journal of Computer Vision</i> , 2005 , 62, 61-81	10.6	32
43	Seeing the Arrow of Time 2014 ,		29
42	You Said That?: Synthesising Talking Faces from Audio. <i>International Journal of Computer Vision</i> , 2019 , 127, 1767-1779	10.6	28
41	Learnable PINs: Cross-modal Embeddings for Person Identity. <i>Lecture Notes in Computer Science</i> , 2018 , 73-89	0.9	27
40	On-the-fly learning for visual search of large-scale image and video datasets. <i>International Journal of Multimedia Information Retrieval</i> , 2015 , 4, 75-93	2.4	25
39	Learning to Detect Partially Overlapping Instances 2013 ,		24
38	2020 ,		23
37	Learning equivariant structured output SVM regressors 2011 ,		21

36	Total Cluster 2014 ,		20
35	Time-lapse imagery and volunteer classifications from the Zooniverse Penguin Watch project. <i>Scientific Data</i> , 2018 , 5, 180124	8.2	20
34	On-the-fly specific person retrieval 2012 ,		18
33	Self-supervised Learning of Audio-Visual Objects from Video. <i>Lecture Notes in Computer Science</i> , 2020 , 208-224	0.9	18
32	Smooth-AP: Smoothing the Path Towards Large-Scale Image Retrieval. <i>Lecture Notes in Computer Science</i> , 2020 , 677-694	0.9	18
31	Memory-Augmented Dense Predictive Coding for Video Representation Learning. <i>Lecture Notes in Computer Science</i> , 2020 , 312-329	0.9	18
30	Humanising GrabCut: Learning to segment humans using the Kinect 2011 ,		17
29	2019 ,		15
28	From Benedict Cumberbatch to Sherlock Holmes: Character Identification in TV series without a Script 2017 ,		14
27	BiCoS: A Bi-level co-segmentation method for image classification 2011 ,		13
26	Amplifying Key Cues for Human-Object-Interaction Detection. <i>Lecture Notes in Computer Science</i> , 2020 , 248-265	0.9	12
25	3D Shape Attributes 2016 ,		10
24	Synthetic Humans for Action Recognition from Unseen Viewpoints. <i>International Journal of Computer Vision</i> , 2021 , 129, 2264-2287	10.6	10
23	TeachText: CrossModal Generalized Distillation for Text-Video Retrieval 2021 ,		9
22	Automated audiovisual behavior recognition in wild primates. <i>Science Advances</i> , 2021 , 7, eabi4883	14.3	8
21	Self-Supervised Learning of Class Embeddings from Video 2019 ,		8
20	Learning to Read by Spelling 2018 ,		6
19	Semi-Supervised Learning with Scarce Annotations 2020 ,		6

18	Learning to Predict 3D Surfaces of Sculptures from Single and Multiple Views. <i>International Journal of Computer Vision</i> , 2019 , 127, 1780-1800	10.6	6
17	Identifying Scoliosis in Population-Based Cohorts: Automation of a Validated Method Based on Total Body Dual Energy X-ray Absorptiometry Scans. <i>Calcified Tissue International</i> , 2020 , 106, 378-385	3.9	5
16	Immediate, Scalable Object Category Detection 2014 ,		5
15	Slow-Fast Auditory Streams for Audio Recognition 2021 ,		5
14	Human pose search using deep poselets 2015 ,		4
13	From Same Photo: Cheating on Visual Kinship Challenges. <i>Lecture Notes in Computer Science</i> , 2019 , 654-668	6.9	4
12	LAEO-Net++: revisiting people Looking At Each Other in videos. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 2021 , PP,	13.3	4
11	Sight to Sound: An End-to-End Approach for Visual Piano Transcription 2020 ,		3
10	Discovery of Rare Phenotypes in Cellular Images Using Weakly Supervised Deep Learning 2017 ,		3
9	Processing citizen science- and machine-annotated time-lapse imagery for biologically meaningful metrics. <i>Scientific Data</i> , 2020 , 7, 102	8.2	2
8	LSD-C: Linearly Separable Deep Clusters 2021 ,		2
7	The Use and Reuse of Printed Illustrations in 15th-Century Venetian Editions 2020 ,		2
6	Scaling Up Sign Spotting Through Sign Language Dictionaries. <i>International Journal of Computer Vision</i> , ¹	10.6	2
5	Where Is the Archive?. <i>Photography and Culture</i> , 2010 , 3, 331-342	0.1	1
4	2019 ,		1
3	Automated Video Face Labelling for Films and TV Material. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 2020 , 42, 780-792	13.3	0
2	Illuminance Flow Estimation by Regression. <i>International Journal of Computer Vision</i> , 2010 , 90, 304-312	10.6	
1	Bounding an archiving: assessing the relative completeness of the Jacques Toussele archive using pattern-matching and face-recognition. <i>Visual Studies</i> , ¹⁻²⁵	0.3	

