

Li Fei-Fei

List of Publications by Citations

Source: <https://exaly.com/author-pdf/12097401/li-fei-fei-publications-by-citations.pdf>

Version: 2024-04-24

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.

95
papers

47,215
citations

49
h-index

97
g-index

97
ext. papers

64,830
ext. citations

8
avg, IF

7.86
L-index

#	Paper	IF	Citations
95	ImageNet Large Scale Visual Recognition Challenge. <i>International Journal of Computer Vision</i> , 2015 , 115, 211-252	10.6	14539
94	ImageNet: A large-scale hierarchical image database 2009 ,		14481
93	Large-Scale Video Classification with Convolutional Neural Networks 2014 ,		2868
92	Perceptual Losses for Real-Time Style Transfer and Super-Resolution. <i>Lecture Notes in Computer Science</i> , 2016 , 694-711	0.9	1999
91	Deep visual-semantic alignments for generating image descriptions 2015 ,		1529
90	One-shot learning of object categories. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 2006 , 28, 594-611	13.3	1259
89	Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. <i>Computer Vision and Image Understanding</i> , 2007 , 106, 59-70	4.3	1141
88	Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. <i>International Journal of Computer Vision</i> , 2017 , 123, 32-73	10.6	1083
87	Unsupervised Learning of Human Action Categories Using Spatial-Temporal Words. <i>International Journal of Computer Vision</i> , 2008 , 79, 299-318	10.6	961
86	What, where and who? Classifying events by scene and object recognition 2007 ,		406
85	Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories		386
84	Deep Visual-Semantic Alignments for Generating Image Descriptions. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 2017 , 39, 664-676	13.3	341
83	Modeling mutual context of object and human pose in human-object interaction activities 2010 ,		322
82	CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning 2017 ,		321
81	Auto-DeepLab: Hierarchical Neural Architecture Search for Semantic Image Segmentation 2019 ,		301
80	Image retrieval using scene graphs 2015 ,		292
79	Natural scene categories revealed in distributed patterns of activity in the human brain. <i>Journal of Neuroscience</i> , 2009 , 29, 10573-81	6.6	253

78	Neural mechanisms of rapid natural scene categorization in human visual cortex. <i>Nature</i> , 2009 , 460, 94-750.4	225
77	What do we perceive in a glance of a real-world scene?. <i>Journal of Vision</i> , 2007 , 7, 10	0.4 225
76	Modeling Temporal Structure of Decomposable Motion Segments for Activity Classification. <i>Lecture Notes in Computer Science</i> , 2010 , 392-405	0.9 219
75	Visual Relationship Detection with Language Priors. <i>Lecture Notes in Computer Science</i> , 2016 , 852-869	0.9 218
74	ImageNet: A large-scale hierarchical image database 2009 ,	208
73	Learning latent temporal structure for complex event detection 2012 ,	187
72	2010 ,	162
71	Simple line drawings suffice for functional MRI decoding of natural scene categories. <i>Proceedings of the National Academy of Sciences of the United States of America</i> , 2011 , 108, 9661-6	11.5 141
70	Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 2012 , 34, 1691-703	13.3 134
69	What Does Classifying More Than 10,000 Image Categories Tell Us?. <i>Lecture Notes in Computer Science</i> , 2010 , 71-84	0.9 134
68	Inferring and Executing Programs for Visual Reasoning 2017 ,	133
67	Towards total scene understanding: Classification, annotation and segmentation in an automatic framework 2009 ,	130
66	What's the Point: Semantic Segmentation with Point Supervision. <i>Lecture Notes in Computer Science</i> , 2016 , 549-565	0.9 130
65	OPTIMOL: Automatic Online Picture Collection via Incremental Model Learning. <i>International Journal of Computer Vision</i> , 2010 , 88, 147-168	10.6 121
64	Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos. <i>International Journal of Computer Vision</i> , 2018 , 126, 375-389	10.6 120
63	Differential connectivity within the Parahippocampal Place Area. <i>NeuroImage</i> , 2013 , 75, 228-237	7.9 102
62	Tool Detection and Operative Skill Assessment in Surgical Videos Using Region-Based Convolutional Neural Networks 2018 ,	99
61	Spatial-Temporal correlatons for unsupervised action classification 2008 ,	97

60	The Unreasonable Effectiveness of Noisy Data for Fine-Grained Recognition. <i>Lecture Notes in Computer Science</i> , 2016 , 301-320	0.9	93
59	Object Bank: An Object-Level Image Representation for High-Level Visual Recognition. <i>International Journal of Computer Vision</i> , 2014 , 107, 20-39	10.6	92
58	Generating Semantically Precise Scene Graphs from Textual Descriptions for Improved Image Retrieval 2015 ,		92
57	Why does natural scene categorization require little attention? Exploring attentional requirements for natural and synthetic stimuli. <i>Visual Cognition</i> , 2005 , 12, 893-924	1.8	86
56	OPTIMOL: automatic Online Picture collecTION via Incremental MOdel Learning 2007 ,		73
55	Bedside Computer Vision - Moving Artificial Intelligence from Driver Assistance to Patient Safety. <i>New England Journal of Medicine</i> , 2018 , 378, 1271-1273	59.2	66
54	Fine-Grained Recognition in the Wild: A Multi-task Domain Adaptation Approach 2017 ,		66
53	Two Distinct Scene-Processing Networks Connecting Vision and Memory. <i>ENeuro</i> , 2016 , 3,	3.9	66
52	Scalable multi-label annotation 2014 ,		64
51	Love Thy Neighbors: Image Annotation by Exploiting Image Metadata 2015 ,		61
50	Reasoning about Object Affordances in a Knowledge Base Representation. <i>Lecture Notes in Computer Science</i> , 2014 , 408-424	0.9	61
49	Distinct contributions of functional and deep neural network features to representational similarity of scenes in human brain and behavior. <i>ELife</i> , 2018 , 7,	8.9	60
48	Learning semantic relationships for better action retrieval in images 2015 ,		59
47	Illuminating the dark spaces of healthcare with ambient intelligence. <i>Nature</i> , 2020 , 585, 193-202	50.4	52
46	Connectionist Temporal Modeling for Weakly Supervised Action Labeling. <i>Lecture Notes in Computer Science</i> , 2016 , 137-153	0.9	49
45	Towards fairer datasets 2020 ,		49
44	Towards Scalable Dataset Construction: An Active Learning Approach. <i>Lecture Notes in Computer Science</i> , 2008 , 86-98	0.9	45
43	. <i>Proceedings of the IEEE</i> , 2010 , 98, 1453-1466	14.3	44

42	A computer vision system for deep learning-based detection of patient mobilization activities in the ICU. <i>Npj Digital Medicine</i> , 2019 , 2, 11	15.7	41
41	Using Dependent Regions for Object Categorization in a Generative Framework		40
40	What you see is what you expect: rapid scene understanding benefits from prior experience. <i>Attention, Perception, and Psychophysics</i> , 2015 , 77, 1239-51	2	38
39	Action Genome: Actions As Compositions of Spatio-Temporal Scene Graphs 2020 ,		38
38	Visual scenes are categorized by function. <i>Journal of Experimental Psychology: General</i> , 2016 , 145, 82-94	4.7	37
37	Visual categorization is automatic and obligatory: evidence from Stroop-like paradigm. <i>Journal of Vision</i> , 2014 , 14,	0.4	37
36	Parcellating connectivity in spatial maps. <i>PeerJ</i> , 2015 , 3, e784	3.1	37
35	Detecting Avocados to Zucchini: What Have We Done, and Where Are We Going? 2013 ,		36
34	Embracing Error to Enable Rapid Crowdsourcing 2016 ,		35
33	Online Developmental Science to Foster Innovation, Access, and Impact. <i>Trends in Cognitive Sciences</i> , 2020 , 24, 675-678	14	31
32	Basic level category structure emerges gradually across human ventral visual cortex. <i>Journal of Cognitive Neuroscience</i> , 2015 , 27, 1427-46	3.1	27
31	Human-Object Interactions Are More than the Sum of Their Parts. <i>Cerebral Cortex</i> , 2017 , 27, 2276-2288	5.1	25
30	Voxel-level functional connectivity using spatial regularization. <i>NeuroImage</i> , 2012 , 63, 1099-106	7.9	24
29	Binding is a local problem for natural objects and scenes. <i>Vision Research</i> , 2005 , 45, 3133-44	2.1	24
28	Learning task-oriented grasping for tool manipulation from simulated self-supervision. <i>International Journal of Robotics Research</i> , 2020 , 39, 202-216	5.7	23
27	Discovering Object Functionality 2013 ,		22
26	Typicality sharpens category representations in object-selective cortex. <i>NeuroImage</i> , 2016 , 134, 170-179	7.9	22
25	Automated abnormality detection in lower extremity radiographs using deep learning. <i>Nature Machine Intelligence</i> , 2019 , 1, 578-583	22.5	22

24	Good exemplars of natural scene categories elicit clearer patterns than bad exemplars but not greater BOLD activity. <i>PLoS ONE</i> , 2013 , 8, e58594	3.7	19
23	Learning Task-Oriented Grasping for Tool Manipulation from Simulated Self-Supervision		19
22	Pinpointing the peripheral bias in neural scene-processing networks during natural viewing. <i>Journal of Vision</i> , 2016 , 16, 9	0.4	17
21	Situational Fusion of Visual Representation for Visual Navigation 2019 ,		15
20	Evidence for similar patterns of neural activity elicited by picture- and word-based representations of natural scenes. <i>NeuroImage</i> , 2017 , 155, 422-436	7.9	13
19	Assessing the accuracy of automatic speech recognition for psychotherapy. <i>Npj Digital Medicine</i> , 2020 , 3, 82	15.7	13
18	Task-set switching with natural scenes: measuring the cost of deploying top-down attention. <i>Journal of Vision</i> , 2007 , 7, 9.1-12	0.4	13
17	RubiksNet: Learnable 3D-Shift for Efficient Video Action Recognition. <i>Lecture Notes in Computer Science</i> , 2020 , 505-521	0.9	13
16	Extracting Moving People from Internet Videos. <i>Lecture Notes in Computer Science</i> , 2008 , 527-540	0.9	12
15	Ethical issues in using ambient intelligence in health-care settings. <i>The Lancet Digital Health</i> , 2021 , 3, e115-e123	14.4	11
14	A multi-view probabilistic model for 3D object classes 2009 ,		9
13	Automatic detection of hand hygiene using computer vision technology. <i>Journal of the American Medical Informatics Association: JAMIA</i> , 2020 , 27, 1316-1320	8.6	9
12	Categorization influences detection: A perceptual advantage for representative exemplars of natural scene categories. <i>Journal of Vision</i> , 2017 , 17, 21	0.4	8
11	Vision-based Estimation of MDS-UPDRS Gait Scores for Assessing Parkinson's Disease Motor Severity. <i>Lecture Notes in Computer Science</i> , 2020 , 12263, 637-647	0.9	8
10	Quantifying Parkinson's disease motor severity under uncertainty using MDS-UPDRS videos. <i>Medical Image Analysis</i> , 2021 , 73, 102179	15.4	7
9	Two distinct scene processing networks connecting vision and memory		6
8	Audio-Visual Speaker Localization Using Graphical Models 2006 ,		5
7	Scene Graph Prediction with Limited Labels. <i>Proceedings of the IEEE International Conference on Computer Vision</i> , 2019 , 2019, 2580-2590	3.3	4

6	A multi-view probabilistic model for 3D object classes		3
5	Towards total scene understanding: Classification, annotation and segmentation in an automatic framework		2
4	Visual Intelligence through Human Interaction. <i>Human-computer Interaction Series</i> , 2021 , 257-314	0.6	2
3	Deep Affordance Foresight: Planning Through What Can Be Done in the Future 2021 ,		2
2	Neural Event Semantics for Grounded Language Understanding. <i>Transactions of the Association for Computational Linguistics</i> , 2021 , 9, 875-890	5.6	1
1	Thoracic Disease Identification and Localization with Limited Supervision. <i>Advances in Computer Vision and Pattern Recognition</i> , 2019 , 139-161	1.1	