

Fei-Fei Li

(Publish as L. Fei-Fei)

353 Serra Mall, Room 2A-246 | Stanford, CA 94305
+1-650-725-3860 | feifeili@cs.stanford.edu | svl.stanford.edu
(Last update: 2018.10)

EDUCATION

- CALIFORNIA INSTITUTE OF TECHNOLOGY** **Pasadena, CA, U.S.A.**
Doctor of Philosophy in Electrical Engineering (Ph.D) **2001 - 2005**
- Advisors: Pietro Perona (primary) and Christof Koch (secondary)
 - Dissertation: “*Visual Recognition: Computational Models and Human Psychophysics*”
- CALIFORNIA INSTITUTE OF TECHNOLOGY** **Pasadena, CA, U.S.A.**
Master of Science in Electrical Engineering **2001 - 2003**
- Advisors: Pietro Perona (primary) and Christof Koch (secondary)
- PRINCETON UNIVERSITY** **Princeton, NJ, U.S.A.**
Bachelor of Arts in Physics **1995 - 1999**
- Graduated with High Honors
 - Certificates (equivalent of Minor) in Applied & Computational Mathematics, and Engineering Physics

WORK EXPERIENCES - ACADEMIA

- STANFORD UNIVERSITY** **Stanford, CA, U.S.A.**
Professor, Computer Science Department **2018.01 - Present**
Co-Director, Stanford Human-Centered AI Institute (HAI) **2018.10 - Present**
Director, Stanford Artificial Intelligence Lab (SAIL) **2013 - 2018.10**
Associate Professor, Computer Science Department **2012.08 - 2017.12**
Assistant Professor, Computer Science Department **2009.06 - 2012.08**
- Diversity committee, Computer Science Department (2016, 2017)
 - Co-Founder/Co-Director, Stanford AI Lab Outreach Summer “*SAILORS*” (2015 - present)
 - Faculty Advisor, Stanford AI Lab AI women’s group (2015 - present)
 - Founding Director, Stanford-Toyota Human-Centric AI Research Center (09/2015 - 12/2016)
- PRINCETON UNIVERSITY** **Princeton, NJ, U.S.A.**
Assistant Professor, Computer Science Department **2007.01 - 2009.06**
Associated Assistant Professor, Psychology Department **2007.01 - 2009.06**
- UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN** **Urbana-Champaign, IL, U.S.A.**
Assistant Professor, Electrical and Computer Engineering Department **2005.08 - 2006.12**
Associated Assistant Professor, Psychology Department **2005.08 - 2006.12**
- MICROSOFT RESEARCH CENTER CAMBRIDGE** **Cambridge, UK**
Visiting Research Scientist **2005.03 - 2005.06**

WORK EXPERIENCES -- INDUSTRY, NON-PROFIT & MISCELLANEOUS

- AI4ALL - Non-profit organization for AI education and diversity** **Oakland, U.S.A.**
Co-Founder & Chairperson of the Board **2017.03 - present**
- Co-Founder & Director of precursor program -- SAILORS (Stanford AI Lab Outreach Summer program) (2015 - 2017)

GOOGLE INC.**Mountain View, CA, U.S.A.**Chief Scientist of AI/ML, Vice President, Google Cloud AI
(on academic leave from Stanford)**2017.01 - 2018.09**

- Co-founder of Cloud AI business unit
- Overall responsibility of product engineering, product management, basic science research and thought leadership
- Grow the Cloud AI team, business, and partnership
- Chief architect of major Google Cloud AI products such as AutoML Vision, NLP, Translation
- Business leader responsible for acquisition of Kaggle.com
- Business leader responsible for establishing Google AI's China Center in Beijing

ANDREESSON & HOROWITZ INC.**2016.05 - 2017.01**

Professor in Residence

SNAP INC.**2014.11 - 2017.01**

Advisor

VICARIOUS, INC.**2014.05 - present**

Advisor

GAUSSSURGICAL, INC.**2012.02 - present**

Advisor (Technical Board)

ONE H.E.A.R.T FOUNDATION**2002-2004**

Research Scientist

Project: ethnographic research on maternal health in rural and nomadic Tibet

TIBETAN MEDICINE HOSPITAL (Mentsikhang)**1999 - 2000**

Principal Investigator

Project: clinical study on Tibetan medicine treatment for Helicobacter pylori infection in rural and nomadic Tibetan population

HONORS & DISTINCTIONS

- 2017 WITI@UC Athena Award for Academic Leadership, University of California
- 2017 One of Seven Women in Technology honorees, Elle Magazine
- 2016 J.K. Aggarwal Prize, International Association for Pattern Recognition (IAPR)
- 2016 One of the 40 "The great immigrants," Carnegie Foundation
- 2016 IEEE PAMI Mark Everingham Prize
- 2016 Pioneer in AI Research Award, NVidia
- 2015 One of the Leading Global Thinkers of 2015, Foreign Policy
- 2014 IBM Faculty Fellowship Award
- 2012 W.M. Keck Foundation Faculty Scholar
- 2012 Yahoo! Labs Faculty Research Engagement Program (FREP) Award, Yahoo!
- 2012/11 1st Place in PASCAL VOC Action Classification Challenge (internationally recognized computer vision competition)
- 2011 Alfred P. Sloan Fellowship (highly prestigious fellowship awarded to "best scholars in [the current] generation")
- 2010 Best Paper Honorable Mention, IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- 2010 Google Research Award
- 2009 Stanford Terman Fellowship
- 2009 NSF CAREER Award
- 2008 Google Research Award

- 2007 1st Place in Semantic Robot Vision Challenge Software League - NSF / AAI sponsored visual recognition competition
- 2006 Microsoft Research New Faculty Fellowship - highly selective, awarded to “the best new professors in computing disciplines today”
- 2005 IEEE ICCV Best Short Course Prize (with R. Fergus and A. Torralba)
- 1999 - 2002 National Science Foundation Postgraduate Fellowship
- 1999 - 2002 Paul and Daisy Soros Fellowship for New Americans
- 1999 - 2000 Princeton University Martin Dale '53 Fellowship
- 1999 Princeton University Kusaka Memorial Prize in Physics

SPEECHES & SEMINARS

“Artificial Intelligence: Great Power Comes with Great Responsibilities”

2018.06 Expert witness, Congressional Hearing

“Artificial Intelligence: A Deeply Human Pursuit”

2018.04 Invited Lecture for the Lorna Casselton Memorial Lecture Series, Oxford University, Oxford, U.K.

2017.10 Keynote, Grace Hopper Conference, FL, USA

“Towards Ambient Intelligence in AI-assisted Healthcare Spaces”

2018.04 Invited Lecture for Alan Turing Institute, London, U.K.

“Visual Intelligence: Beyond ImageNet”

2017.10 Plenary, International Conference on Intelligent Robots and Systems (IROS), Vancouver, Canada

2017.10 Keynote, Chinese National Computer Congress (CNCC), Fuzhou, China

“ImageNet”

2017.09 Talk, ACM Learning Webinar

2017.05 Keynote, IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, USA

“Visual Intelligence: Shining Light on the Digital World”

2017.05 Inclusive AI, Berkeley, USA

2017.06 Plenary Talk: United Nations & ITU AI for Good Conference, Geneva, Switzerland

“Seeing is Understanding”

2017.03 Invited Talk: AAI Workshop on Visual Understanding, Stanford, USA

“Teaching Computers To See”

2015.03 Invited Talk: TED annual conference 2015, Vancouver, Canada

“Guardian Angels: Towards AI-assisted Care”

2016.05 Invited Talk: Big Data in Biomedicine, Stanford, CA, USA

“Quest for Visual Intelligence”

2017.02 Keynote Talk: Women in Data Science (WiDS), Stanford, U.S.A.

2017.01 Keynote Talk: Future Forum Annual Conference, Beijing, China

2016.12 Plenary Talk: International Conference of Pattern Recognition (ICPR), Cancun, Mexico

2016.11 Distinguished Speaker Seminar: UC Berkeley EECS Dept, U.S.A.

2016.06 Keynote Talk: Stanford-White House Office of Science and Technology Policies (OSTP) “The Future of Artificial Intelligence” Symposium, Stanford, CA, USA

2016.06 Keynote Talk: International Conference on Machine Learning (ICML), New York, USA

2016.03 Keynote Talk: Sandisk Tech Day, CA, USA

2015.11 Invited Talk: Stanford Women in Data Science (WiDS)
2015.12 Keynote Talk: Yahoo! Tech Day, CA, USA
2015.08 Keynote Talk: VM World Conference, CA, USA
2015.06 Plenary Talk: NAACL, CO, USA
2015.05 Keynote Talk: Machine Learning Conference, Amazon, Seattle
2015.04 Keynote Talk: IEEE Big Data Symposium, San Mateo, USA
2015.02 Featured Talk: Strata+Hadoop World 2015
2014.11 Invited Talk: Cognitive Systems Colloquium, Almedan, U.S.
2014.11 Snapchat, USA
2014.09 Keynote Address: British Machine Vision Conference (BMVC), Nottingham, U.K.
2014.08 Plenary Speech: International Conference on Pattern Recognition (ICPR), Stockholm, Sweden
2014.07 Tencent, China
2014.04 Distinguished Lecture: Sonoma State University, Computer Science Department, CA, U.S.
2014.03 Distinguished Lecture: Simon-Fraser University, Canada
2014.01 Adobe, USA
2014.01 Qualcomm, USA
2013.12 General Electric, USA
2013.11 Yahoo! Big Thinker Lecture

“A Tale of Two Senses: Recognizing Pictures and Grounding Words”

2014.09 Keynote Address: Computer Vision + Ontology Applied Cross-disciplinary Technologies (CONTACT) workshop at IEEE European Conference on Computer Vision (ECCV), Zurich, Switzerland

“Human Behavior Understanding”

2014.09 Invited Talk: Human Behavior Understanding workshop at IEEE European Conference on Computer Vision (ECCV), Zurich, Switzerland

“Let’s Reason About Object Affordance”

2014.09 Keynote Address: Visual Perception of Affordances and Functional Visual Primitives for Scene Analysis workshop at IEEE European Conference on Computer Vision (ECCV), Zurich, Switzerland

“Large-Scale Visual Recognition (Powered by Big Data and Big Crowd)”

2013.09 Keynote Address: International Conference on Image Analysis and Processing (ICIAP), Naples, Italy
2013.07 Microsoft Faculty Summit, Redmond, U.S.A.
2013.04 Microsoft Machine Learning Summit, Paris, France

“Fine-Grained Recognition: From Machines to Machine-Crowd Collaboration”

2013.06 Invited Speech: Fine-Grained Visual Categorization workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Portland, U.S.A.

“Understanding Human Activities”

2013.06 Invited Speech: Scene Understanding Workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Portland, U.S.A.

“Recognizing Millions of Categories, and Never Making a Mistake”

2013.06 Keynote Address: Vision and Language Workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Portland, U.S.A.

“Crowds in the Cloud: the Artificial Artificial Technology in Visual Recognition”

2012.10 Invited Speech: Computer Vision for the Web Workshop at European Conference on Computer Vision (ECCV), Florence, Italy

2011.06 Keynote Address: Opus Venture Annual Conference, U.S.A.

“Modeling Mutual Context of Object and Human Pose in Human-Object Interaction Activities”

2011.06 Invited Speech: International Conference on Machine Learning (ICML), Seattle, U.S.A.

“Object Bank: A High-Level Image Representation for Complex Scene Understanding”

2011.06 Invited Speech: Workshop on Learning Architectures, Representations and Optimization for Speech and Visual Information Processing at ICML, Seattle, U.S.A.

“Large-Scale Image Classification: ImageNet and ObjectBank”

2011.06 Microsoft Bing Tech Talk, Redmond, U.S.A.

2011.03 Google Tech Talk, Mountain View, U.S.A.

“Recognizing Human-Object Interaction Activities”

2011.06 Keynote Address: Workshop on Activity Recognition Challenges at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Colorado Springs, U.S.A.

“Building the Forest: Large-Scale Data and Modeling in Computer Vision”

2011.06 Invited Speech: Workshop on Large Scale Learning for Vision at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Colorado Springs, U.S.A.

“Combining Randomization and Discrimination for Fine-Grained Categorization”

2011.06 Invited Speech: Workshop on Fine-Grained Visual Recognition at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Colorado Springs, U.S.A.

“High-Level Visual Recognition (Handling Hidden Structure, High Dimensionality and Large-Scale Data)”

2011.04 Invited Speech, Information Science and Technology Seminar Series at California Institute of Technology, U.S.A.

2010.12 Invited Speech, Princeton University, U.S.A.

2010.02 Invited Speech, Workshop on Machine Learning for Next Generation Computer Vision Challenge at NIPS, Vancouver, Canada

“Visual Recognition Beyond Simple Actions and Isolated Actors and Objects”

2010.06 Keynote Address, 2nd Workshop on User of Context in Video Processing (UCVP) at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), San Francisco, U.S.A.

“Building ImageNet: Keeping Humans in the Loop”

Keynote Address, Workshop on Advancing Computer Vision with Humans in the Loop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), San Francisco, U.S.A.

“What, Where and Whom: What Do Humans See in a Glance of a Scene? And What Can Computers See?”

2010.04 Invited Talk, NSF Workshop on Hybrid Neuro-Computer Vision System, Columbia University, U.S.A.

“Telling Stories in Images: Modeling Hierarchies In and Across Images” (“Towards Total Scene Understanding”, “Telling the Story of an Image”)

2010.03 Invited Talk, Intelligence Seminar, Carnegie Mellon University, U.S.A.

2010.01 Invited Talk, National Taiwan University, Taiwan

2010.01 Invited Talk, National ChiaoTung University, Taiwan

2008.07 Invited Talk, Google Research Lab, China

2008.07 Invited Talk, Chinese Academy of Science, China

2008.05 Keynote Speech, 2nd Annual Perceptual Science Forum, Rutgers University, U.S.A.

2008.03 Invited Talk, University of Maryland, U.S.A.
2007.12 Invited Talk, ONR MURI Workshop, Caltech, U.S.A.

“ImageNet: Crowdsourcing, Benchmarking and Other Cool Things”

2010.03 Invited Talk, VASC Seminar, Carnegie Mellon University, U.S.A.
2009.06 Invited Talk, 2nd Workshop on Internet Vision at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, U.S.A.

“From Bag-of-Words to Total Scene Understanding: Evolution of Topic Models in Visual Recognition”

2009.12 Invited Talk, Workshop on Applications for Topic Models: Text and Beyond at NIPS, Vancouver, U.S.A.

“Natural Scene Categorization: Behaviors, Brains and Computers”

2009.06 Invited Talk, 1st Workshop on Visual Place Categorization at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, U.S.A.

“Understanding Human Motion”

2008.12 Invited Talk, New York University, U.S.A.

“Discovering Meaning in the Visual World”

2007.08 Invited Talk, Microsoft Research Asia (MSRA), Beijing, China
2007.06 Invited Talk, Google Research, NYC, U.S.A.
2007.05 Invited Talk, MIT Machine Vision Colloquium, U.S.A.
2007.02 Invited Talk, COSYNE Workshop, Utah, U.S.A.

“Natural Scene Classification Using Distributed fMRI Activity”

2007.02 Invited Talk, Scene Understanding Symposium, MIT, U.S.A.

“Unsupervised Learning of Human Action Categories Using Spatial-Temporal Words”

2006.09 Invited Talk, CALD Seminar Series, Carnegie Mellon University, U.S.A.

“Natural Scene Categorization in Humans and Computers”

2006.02 Invited Talk, Scene Understanding Symposium, MIT, U.S.A.
2006.01 Invited Talk, AIVRH Seminar Series, UIUC, U.S.A.
2005.12 Invited Talk, Graphics Lab, Princeton University, U.S.A.
2005.12 Invited Talk, Sarnoff Corporation, U.S.A.
2005.12 Invited Talk, Siemens Corporation, U.S.A.
2005.04 Invited Talk, Visual Geometry Group, Oxford University, U.K.

“A Bayesian Framework for Unsupervised One-Shot Learning of Object Categories”

2004.02 Invited Talk, SHAPE Seminar Series, Brown University, U.S.A.
2004.02 Invited Talk, Electrical Engineering Dept, Princeton University, U.S.A.
2003.12 Invited Talk, Computer Science Dept, New York University, U.S.A.
2003.10 Invited Talk, Louvain University, Belgium

“Natural Scene Categorization Without Attention”

2003.10 Invited Talk, Katholieke Universiteit Leuven, Belgium
2003.05 Invited Talk, Statistics Dept, UCLA, U.S.A.
2003.02 Invited Talk, Brain and Cognitive Science Dept, MIT, U.S.A.
2003.01 Invited Talk, Plymouth University, U.K.
2002.12 Invited Talk, CNRS Toulouse, France

PUBLICATIONS

BOOK CHAPTERS, REFEREED JOURNALS AND CONFERENCE PAPERS

([Google Scholar link](#))

1. Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li Fei-Fei, Juan Carlos Niebles. Learning to Decompose and Disentangle Representations for Video Prediction. *Neural Information Processing Systems (NIPS)*, 2018.
2. Nick Haber*, Damian Mrowca*, Li Fei-Fei, Daniel L K. Yamins. Learning to Play with Intrinsically-Motivated Self-Aware Agents. *Neural Information Processing Systems (NIPS)*, 2018.
3. Damian Mrowca*, Chengxu Zhuang*, Elias Wang*, Nick Haber, Li Fei-Fei, Joshua B Tenenbaum, Daniel L. K. Yamins. Flexible Neural Representation for Physics Prediction. *Neural Information Processing Systems (NIPS)*, 2018.
4. Linxi Fan*, Yuke Zhu*, Jiren Zhu, Zihua Liu, Orien Zeng, Anchit Gupta, Joan Creus-Costa, Silvio Savarese, Li Fei-Fei. SURREAL: Open-Source Reinforcement Learning Framework and Robot Manipulation Benchmark. *Conference on Robot Learning (CoRL)*. 2018.
5. Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Anchit Gupta, Emre Orbay, Silvio Savarese, Li Fei-Fei. RoboTurk: A Crowdsourcing Platform for Robotic Skill Learning through Imitation. *Conference on Robot Learning (CoRL)*. 2018.
6. M. Guo, E. Chou, D.-A. Huang, S. Song, S. Yeung, & L. Fei-Fei. NEURAL GRAPH MATCHING NETWORKS FOR FEWSHOT 3D ACTION RECOGNITION. *IEEE European Conference on Computer Vision (ECCV)*. 2018.
7. M. Guo, A. Haque, D.-A. Huang, S. Yeung, & L. Fei-Fei. FOCUS ON THE HARD THINGS: DYNAMIC TASK PRIORITIZATION FOR MULTITASK LEARNING. *IEEE European Conference on Computer Vision (ECCV)*. 2018.
8. J. Zhu*, R. Kaplan*, J. Johnson, & L. Fei-Fei, HIDDEN: HIDING DATA WITH DEEP NETWORKS, *IEEE European Conference on Computer Vision (ECCV)*. 2018.
9. B. Liu, S. Yeung, E. Chou, D.-A. Huang, L. Fei-Fei, & J.C. Niebles. TEMPORAL MODULAR NETWORKS FOR RETRIEVING COMPLEX COMPOSITIONAL ACTIVITIES IN VIDEOS. *IEEE European Conference on Computer Vision (ECCV)*. 2018.
10. Z. Luo, J.-T. Hsieh, L. Jiang, J.C. Niebles, & L. Fei-Fei. GRAPH DISTILLATION FOR ACTION DETECTION WITH PRIVILEGED INFORMATION. *IEEE European Conference on Computer Vision (ECCV)*. 2018.
11. N. Haber*, D. Mrowca*, L. Fei-Fei, & D. L. K. Yamins. EMERGENCE OF STRUCTURED BEHAVIORS FROM CURIOSITY-BASED INTRINSIC MOTIVATION. *Annual Meeting of Cognitive Science Society*. 2018.
12. Z. Luo*, J.-T. Hsieh*, N. Balachandar, S. Yeung, G. Pusiol, J. Luxenberg, G. Li, L.-J. Li, N. L. Downing, A. Milstein, & L. Fei-Fei. COMPUTER VISION-BASED DESCRIPTIVE ANALYTICS OF SENIORS' DAILY ACTIVITIES FOR LONG-TERM HEALTH MONITORING. *International Conference on Machine Learning and Healthcare (MLHC)*. 2018.
13. K. Fang, Y. Zhu, A. Garg, A. Kurenkov, V. Mehta, L. Fei-Fei, & S. Savarese. LEARNING TASK-ORIENTED GRASPING FOR TOOL MANIPULATION FROM SIMULATED SELF-SUPERVISION. *Robotics: Science and Systems (RSS)*. 2018.
14. D. Xu*, S. Nair*, Y. Zhu, J. Gao, A. Garg, L. Fei-Fei, & S. Savarese. NEURAL TASK PROGRAMMING: LEARNING TO GENERALIZE ACROSS HIERARCHICAL TASKS. *IEEE International Conference on Robotics and Automation (ICRA)*. 2018.

15. A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, & A. Alahi. SOCIAL GAN: SOCIALLY ACCEPTABLE TRAJECTORIES WITH GENERATIVE ADVERSARIAL NETWORKS. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
16. J. Johnson, A. Gupta, & L. Fei-Fei. IMAGE GENERATION FROM SCENE GRAPHS. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
17. R. Krishna, I. Chami, M. Bernstein, & L. Fei-Fei. REFERRING RELATIONSHIPS. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
18. D.-A. Huang*, S.I Buch*, L. Dery, A. Garg, L. Fei-Fei, & J.C. Niebles. FINDING "IT": WEAKLY-SUPERVISED REFERENCE-AWARE VISUAL GROUNDING IN INSTRUCTIONAL VIDEO. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
19. D.-A. Huang, V. Ramanathan, D. Mahajan, L. Torresani, M. Paluri, L. Fei-Fei, & J.C. Niebles. WHAT MAKES A VIDEO A VIDEO: ANALYZING TEMPORAL INFORMATION IN VIDEO UNDERSTANDING MODELS AND DATASETS. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
20. S. Yeung, L. N. Downing, L. Fei-Fei, & A. Milstein. Computer Vision – From Driver Assistance to Hospital Safety. *New England Journal of Medicine (NEJM)*. 2018.
21. Z. Luo, Y. Zou, J. Hoffman, & L. Fei-Fei. Label Efficient Learning of Transferable Representations across Domains and Tasks. *Neural Information Processing Systems (NIPS)*. 2017.
22. A. Jin, S. Yeung, J. Jopling, J. Krause, D. Azagury, A. Milstein, & Li Fei-Fei. Tool Detection and Operative Skill Assessment in Surgical Videos Using Region-Based Convolutional Neural Networks. *Machine Learning for Healthcare (ML4H) Workshop at NIPS*. 2017. (Best paper award)
23. O. Russakovsky, N. Jin, M. Andriluka, G. Mori, & L. Fei-Fei. Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos. *International Journal of Computer Vision (IJCV)*. 2017.
24. A. Gupta, J. Johnson, A. Alahi, & L. Fei-Fei. Characterizing and Improving Stability in Neural Style Transfer. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2017.
25. T. Gebru, J. Hoffman & L. Fei-Fei. Fine-grained Recognition in the Wild: A Multi-Task Domain Adaptation Approach. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2017.
26. Y. Zhu*, D. Gordon*, E. Kolve, D. Fox, L. Fei-Fei, A. Gupta, R. Mottaghi, & A. Farhadi, Visual Semantic Planning using Deep Successor Representations. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2017.
27. R. Krishna, K. Hata, F. Ren, L. Fei-Fei, J.C. Niebles. Dense-Captioning Events in Videos. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2017.
28. J. Johnson, B. Hariharan, L. van der Maaten, J. Hoffman, L. Fei-Fei, C. L. Zitnick, R. Girshick. Inferring and Executing Programs for Visual Reasoning. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2017.
29. A. Mandlekar*, Y. Zhu*, A. Garg*, L. Fei-Fei, S. Savarese. Adversarially Robust Policy Learning through Active Construction of Physically-Plausible Perturbations. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017.
30. S. Buch, V. Escorcia, B. Ghanem, L. Fei-Fei, J.C. Niebles. End-to-End, Single-Stream Temporal Action Detection in Untrimmed Videos. *The British Machine Vision Conference (BMVC)*. 2017.
31. A. Haque, A. Alahi, M. Guo, S. Yeung, Z. Luo, A. Rege, A. Singh, J. Jopling, L. Downing, W. Beninati, T. Platchek, A. Milstein, & L. Fei-Fei. *International Conference on Machine Learning and Healthcare (MLHC)*. 2017
32. D.-A. Huang, J. J. Lim, L. Fei-Fei, and J.C. Niebles. Unsupervised Visual-Linguistic Reference Resolution in Instructional Videos. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
33. D. Xu, Y. Zhu, C. B. Choy, & L. Fei-Fei. Scene Graph Generation by Iterative Message Passing. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.

34. Y. Zhu, J. J. Lim, & L. Fei-Fei. Knowledge Acquisition for Visual Question Answering via Iterative Querying. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
35. S. Yeung, V. Ramanathan, O. Russakovsky, L. Shen, G. Mori, & L. Fei-Fei. Learning to learn from noisy web videos. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
36. J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C.L. Zitnick, R. Girshick. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
37. J. Krause, J. Johnson, R. Krishna, & L. Fei-Fei. A Hierarchical Approach for Generating Descriptive Image Paragraphs. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
38. Z. Luo, B. Peng, D.-A. Huang, A. Alahi, & L. Fei-Fei. Unsupervised Learning of Long-Term Motion Dynamics for Videos. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
39. K. Nakamura, S. Yeung, A. Alahi, & L. Fei-Fei. Jointly Learning Energy Expenditures and Activities using Egocentric Multimodal Signals. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
40. A. Alahi, V. Ramanathan, & L. Fei-Fei. Tracking Millions of Humans in Crowded Space. Book chapter in *Group and Crowd Behavior for Computer Vision*. 2017
41. A. Alahi, V. Ramanathan, K. Goel, A. Robicquet, A. A. Sadeghian, L. Fei-Fei, & S. Savarese. Learning to Predict Human Behaviour in Crowded Scenes. Book chapter in *Group and Crowd Behavior for Computer Vision*. 2017
42. T. Gebru, J. Krause, Y. Wang, D. Chen, J. Deng, & L. Fei-Fei. Fine-Grained Car Detection for Visual Census Estimation. *AAAI Conference on Artificial Intelligence (AAAI)*. 2017
43. T. Gebru, J. Krause, J. Deng, & L. Fei-Fei. Scalable Annotation of Fine-Grained Objects Without Experts. *ACM Conference on Human Factors in Computing and Systems (CHI)*. 2017.
44. A. Alahi, J. Wilson, L. Fei-Fei, & S. Savarese. Unsupervised Camera Localization in Crowded Spaces. *IEEE International Conference on Robotics and Automation (ICRA)*. 2017
45. Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, & A. Farhadi. Target-driven Visual Navigation in Indoor Scenes using Deep Reinforcement Learning. *IEEE International Conference on Robotics and Automation (ICRA)*. 2017
46. R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L. Jia-Li, D.A. Shamma, M. Bernstein, & L. Fei-Fei. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision (IJCV)*. 2017
47. K. Hata, R. Krishna, L. Fei-Fei, & M. Bernstein. A Glimpse Far into the Future: Understanding Long-term Crowd Worker Accuracy. *ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*. 2017
48. G. Pusiol, A. Esteva, S. S. Hall, M. Frank, A. Milstein, & L. Fei-Fei. Vision-Based Classification of Developmental Disorders Using Eye-Movements. *International Conference on Medical Image Computing & Computer Assisted Intervention (MICCAI)*. 2016.
49. C. Lu*, R. Krishna*, M. Bernstein, & L. Fei-Fei. Visual Relationship Detection with Language Priors. *IEEE European Conference on Computer Vision (ECCV)*. 2016.
50. D.-A. Huang, L. Fei-Fei, J.C. Niebles. Connectionist Temporal Modeling for Weakly Supervised Action Labeling. *IEEE European Conference on Computer Vision (ECCV)*. 2016.
51. A. Haque, B. Peng*, Z. Luo*, A. Alahi, S. Yeung, & L. Fei-Fei. Towards Viewpoint Invariant 3D Human Pose Estimation. *IEEE European Conference on Computer Vision (ECCV)*. 2016.

52. J. Johnson, A. Alahi, & L. Fei-Fei. Perceptual Losses for Real-time Style Transfer and Single Image Super-Resolution. *IEEE European Conference on Computer Vision (ECCV)*. 2016.
53. A. Bearman, O. Russakovsky, V. Ferrari & L. Fei-Fei. What's the point: Semantic segmentation with point supervision. *IEEE European Conference on Computer Vision (ECCV)*. 2016.
54. J. Krause, B. Sapp, A. Howard, H. Zhou, A. Toshev, T. Duerig, J. Philbin, & L. Fei-Fei. The Unreasonable Effectiveness of Noisy Data for Fine-Grained Recognition. *IEEE European Conference on Computer Vision (ECCV)*. 2016.
55. J. Johnson*, A. Karpathy*, & L. Fei-Fei. DenseCap: Fully Convolutional Localization Networks for Dense Captioning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
56. Y. Zhu, O. Groth, M. Bernstein, & L. Fei-Fei. Visual7W: Grounded Question Answering in Images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
57. A. Haque, A. Alahi, & L. Fei-Fei. Recurrent Attention Models for Depth-Based Person Identification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
58. S. Yeung, O. Russakovsky, G. Mori, & L. Fei-Fei. End-to-end Learning of Action Detection from Frame Glimpses in Videos. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
59. A. Alahi*, K. Goel*, V. Ramanathan, A. Robicquet, L. Fei-Fei, & S Savarese. Social LSTM: Human Trajectory Prediction in Crowded Spaces. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
60. V. Ramanathan, J. Huang, S. Abu-El-Haija, A. Gorban, K. Murphy, & L. Fei-Fei. Detecting Events and Key Actors in Multi-person Videos. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
61. M. R. Greene, C. Baldassano, D. M. Beck, & L. Fei-Fei. Visual Scenes Are Categorized by Function. *Journal of Experimental Psychology: General*. 2016.
62. C. Baldassano, D. M. Beck, & L. Fei-Fei. Human-object Interactions are More than the Sum of Their Parts. *Cerebral Cortex*. 2016.
63. R. Krishna, K. Hata, S. Chen, J. Kravitz, D. A. Shamma, L. Fei-Fei, & M. S. Bernstein. Embracing Error to Enable Rapid Crowdsourcing. *CHI: ACM Conference on Human Factors in Computing Systems*. 2016.
64. M. R. Greene, C. Baldassano, D. M. Beck, & L. Fei-Fei. Visual Scenes Are Categorized by Function. *Journal of Experimental Psychology: General*. 2016.
65. M. E. Vachovsky, G. Wu, S. Chaturapruek, O. Russakovsky, R. Sommer, & L. Fei-Fei. Towards More Gender Diversity in CS through an Artificial Intelligence Summer Program for High School Girls. *Special Interest Group on Computer Science Education (SIGCSE)*. 2016.
66. M. R. Greene, A. P. Botros, D. M. Beck, & L. Fei-Fei. What You See is What You Expect: Rapid Scene Understanding Benefits from Prior Experience. *Attention, Perception, & Psychophysics* 2015.
67. S. Schuster, R. Krishna, A. Chang, & L. Fei-Fei, & C. D. Manning. Generating Semantically Precise Scene Graphs from Textual Descriptions for Improved Image Retrieval. *Empirical Methods in Natural Language Processing (EMNLP) - Vision and Language Workshop 2015 (oral)*. 2015.
68. M. C. Iordan, A. Joulin, D. M. Beck, & L. Fei-Fei. Locally-Optimized Inter-Subject Alignment of Functional Cortical Regions. *Machine Learning and Interpretation in Neuroimaging (MLINI) Workshop, NIPS*. 2015.
69. A. Alahi, A. Haque, & L. Fei-Fei. RGB-W: When Vision Meets Wireless. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2015.
70. K. Tang, M. Paluri, L. Fei-Fei, R. Fergus, & L. Bourdev. Improving Image Classification with Location Context. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2015.
71. V. Ramanathan, K. Tang, G. Mori & L. Fei-Fei. Learning temporal embeddings for complex video analysis. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2015.

72. J. Johnson, L. Ballan, & L. Fei-Fei. Love Thy Neighbors: Image Annotation by Exploiting Image Metadata. *IEEE and CVF International Conference on Computer Vision (ICCV)*. 2015.
73. A. Karpathy, & L. Fei-Fei. Deep Visual-Semantic Alignments for Generating Image Descriptions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.
74. J. Johnson, R. Krishna, M. Stark, L-J Li, D. A. Shamma, M. Bernstein, & L. Fei-Fei. Image retrieval using scene graphs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.
75. J. Krause, H. Jin, J. Yang, & L. Fei-Fei. Fine-grained recognition without part annotations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.
76. O. Russakovsky, L-J Li, & L. Fei-Fei. Best of both worlds: human-machine collaboration for dense object annotation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.
77. O. Russakovsky*, J. Deng*, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*. 2015.
78. C. Baldassano, D. M. Beck, and L. Fei-Fei, Parcellating connectivity in spatial maps. *PeerJ*, 2015.
79. M. C. Iordan, M. R. Greene, D. M. Beck, and L. Fei-Fei, Basic Level Category Structure Emerges Gradually Across Human Ventral Visual Cortex. *Journal of Cognitive Neuroscience*, 2015.
80. A. Karpathy, A. Joulin and L. Fei-Fei. Deep Fragment Embeddings for Bidirectional Image-Sentence Mapping. *Neural Information Processing Systems (NIPS)*. 2014.
81. V. Ramanathan, A. Joulin, P. Liang and L. Fei-Fei. Joint Person Naming in Videos and Coreference Resolution in Text. *IEEE European Conference on Computer Vision (ECCV)*. 2014.
82. Y. Zhu, A. Fathi, and L. Fei-Fei. Reasoning About Object Affordance in a Knowledge Base Representation. *IEEE European Conference on Computer Vision (ECCV)*. 2014.
83. A. Joulin, K. Tang, and L. Fei-Fei. Efficient Image and Video Co-localization with Frank-Wolfe Algorithm. *IEEE European Conference on Computer Vision (ECCV)*. 2014.
84. G. Pusiol, L. Soriano, L. Fei-Fei, & M. C. Frank. Discovering the Signatures of Joint Attention in Child-Caregiver Interaction. *CogSci*. 2014.
85. A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, & L. Fei-Fei. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.
86. A. Alahi, V. Ramanathan, & L. Fei-Fei. Socially-aware Large-scale Crowd Forecasting. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.
87. K. Tang, A. Joulin, L-J Li, & L. Fei-Fei. Co-localization in Real-World Images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.
88. J. Deng, O. Russakovsky, J. Krause, M. Bernstein, A. C. Berg and L. Fei-Fei. Scalable Multi-Label Annotation. *ACM CHI Conference on Human Factors in Computing Systems (CHI)*. 2014.
89. M. Greene, and L. Fei-Fei. Visual Categorization is Automatic and Obligatory: Evidence from a Stroop-like Paradigm. *Journal of Vision*. 2014.
90. B. Yao, J. Ma, and L. Fei-Fei. Discovering Object Functionality. *International Conference on Computer Vision (ICCV)*. 2013.
91. O. Russakovsky, J. Deng, Z. Huang, A. Berg, and L. Fei-Fei. Detecting avocados to zucchinis: what have we done, and where are we going? *International Conference on Computer Vision (ICCV)*. 2013.
92. V. Ramanathan, P. Liang, and L. Fei-Fei. Video Event Understanding using Natural Language Descriptions. *International Conference on Computer Vision (ICCV)*. 2013.
93. K. Tang, B. Yao, L. Fei-Fei, and D. Koller. Combining the Right Features for Complex Event Recognition. *International Conference on Computer Vision (ICCV)*. 2013.
94. C. Baldassano, D.M. Beck, and L. Fei-Fei. Differential Connectivity Within the Parahippocampal Place Area. *NeuroImage*. 2013.
95. V. Ramanathan, B. Yao and L. Fei-Fei. Social Role Discovery in Human Events. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2013

96. K. Tang, R. Skthankar, J. Yagnik and L. Fei-Fei. Discriminative Segment Annotation in Weakly Labeled Video. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2013.
97. A. Karpathy, S. Miller, and L. Fei-Fei. Object Discovery in 3D Scenes via Shape Analysis. *International Conference on Robotics and Automation (ICRA)*. 2013.
98. K. Tang, V. Ramanathan, L. Fei-Fei and D. Koller. Shifting Weights: Adapting Object Detectors from Image to Video. *Neural Information Processing Systems (NIPS)*. 2012.
99. O. Russakovsky, Y. Lin, K. Yu, and L. Fei-Fei. Object-centric Spatial Pooling for Image Classification. *European Conference on Computer Vision (ECCV)*. 2012.
100. B. Yao and L. Fei-Fei. Action Recognition with Exemplar Based 2.5D Graph Matching. *European Conference on Computer Vision (ECCV)*. 2012.
101. C. Baldassano, M.C. Jordan, D.M. Beck, and L. Fei-Fei. Voxel-Level Functional Connectivity using Spatial Regularization. *NeuroImage*. 2012.
102. H. Su, J. Deng, and L. Fei-Fei. Crowdsourcing Annotations for Visual Object Detection. *AAAI Human Computation Workshop*. 2012.
103. G. Kim, L. Fei-Fei, and E. Xing. Web Image Prediction Using Multivariate Point Processes. *The 18th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 2012.
104. H. Su, A. Yu, and L. Fei-Fei. Efficient Euclidean Projections onto the Intersection of Norm Balls. *International Conference on Machine Learning (ICML)*. 2012.
105. B. Yao, G. Bradski, and L. Fei-Fei. A Codebook-Free and Annotation-Free Approach for Fine-Grained Image Categorization. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.
106. G. Kim, L. Fei-Fei, and E. Xing. Web Image Prediction Using Multivariate Point Processes. *The 18th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 2012.
107. H. Su, A. Yu, and L. Fei-Fei. Efficient Euclidean Projections onto the Intersection of Norm Balls. *International Conference on Machine Learning (ICML)*. 2012.
108. K. Tang, L. Fei-Fei, and D. Koller. Learning Latent Temporal Structure for Complex Event Detection. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.
109. J. Deng, J. Krause, A. Berg, and L. Fei-Fei. Hedging Your Bets: Optimizing Accuracy-Specificity Trade-offs in Large Scale Visual Recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.
110. B. Yao and L. Fei-Fei. Recognizing Human Actions in Still Images by Modeling the Mutual Context of Objects and Human Poses. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*. 34(9):1691-1703, September 2012.
111. J. Deng, A. Berg, S. Santheesh, and L. Fei-Fei. Fast and Balanced: Efficient Label Tree Learning for Large Scale Object Recognition. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2011
112. B. Zhao, L. Fei-Fei and E.P. Xing. Large-Scale Category Structure Aware Image Categorization. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2011
113. B. Yao, X. Jiang, A. Khosla, A.L. Lin, L.J. Guibas, and L. Fei-Fei. Human Action Recognition by Learning Bases of Action Attributes and Parts. *International Conference on Computer Vision (ICCV)*. 2011.
114. G. Kim, E.P. Xing, L. Fei-Fei, and T. Kanade. Distributed cosegmentation via submodular optimization on anisotropic diffusion. *IEEE International Conference on Computer Vision (ICCV)*. 2011.
115. M. Savva, N. Kong, A. Chhajta, L. Fei-Fei, M. Agrawala, and J. Heer. ReV ReVision: Automated Classification, Analysis and Redesign of Chart Images: Automated Classification, Analysis and Redesign of Chart Images. *ACM Symposium on User Interface Software and Technology (UIST)*. 2011

116. B. Yao, A. Khosla, and L. Fei-Fei. Classifying Actions and Measuring Action Similarity by Modeling the Mutual Context of Objects and Human Poses. *International Conference on Machine Learning (ICML)*. 2011.
117. D. B. Walther, B. Chai, E. Caddigan, D. M. Beck and L. Fei-Fei. Simple line drawings suffice for functional MRI decoding of natural scene categories. *Proc. Nat. Acad. of Sci (PNAS)*. doi: 10.1073/pnas.1015666108, vol. 108 (no. 23): pp9661-9666. 2011.
118. J. Deng, A. Berg, and L. Fei-Fei. Hierarchical Semantic Indexing for Large Scale Image Retrieval. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2011.
119. B. Yao, A. Khosla, and L. Fei-Fei. Combining Randomization and Discrimination for Fine-Grained Image Categorization. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2011.
120. B. Zhao, L. Fei-Fei, and E.P. Xing, Online Detection of Unusual Events in Videos via Dynamic Sparse Coding. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2011.
121. L.-J. Li, H. Su, E.P. Xing and L. Fei-Fei, Object Bank: A High-Level Image Representation for Scene Classification & Semantic Feature Sparsification. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2010.
122. J. Zhu, L.-J. Li, L. Fei-Fei and E.P. Xing, Large Margin Learning of Upstream Scene Understanding Models. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2010.
123. J. C. Niebles, C.-W. Chen and L. Fei-Fei. Modeling Temporal Structure of Decomposable Motion Segments for Activity Classification. *Proceedings of the 12th European Conference of Computer Vision (ECCV)*. 2010.
124. O. Russakovsky and L. Fei-Fei. Attribute Learning in Large-Scale Datasets. *European Conference of Computer Vision (ECCV) Workshop on Parts and Attributes*. 2010.
125. B. Zhao, L. Fei-Fei and E. P. Xing. Image Segmentation with Topic Random Fields. *Proceedings of the 12th European Conference of Computer Vision (ECCV)*. 2010.
126. J. Deng, A. Berg, K. Li and L. Fei-Fei. What does classifying more than 10,000 image categories tell us? *Proceedings of the 12th European Conference of Computer Vision (ECCV)*. 2010.
127. L. Fei-Fei and L.-J. Li. What, Where and Who? Telling the Story of an Image by Activity Classification, Scene Recognition and Object Categorization. *Book chapter in "Studies in Computational Intelligence- Computer Vision"*. Volume 285/2010, Springer, 2010.
128. S. Savarese, and L. Fei-Fei. Multi-view Object Categorization and Pose Estimation. *Book chapter in "Studies in Computational Intelligence- Computer Vision"*. Volume 285/2010, Springer, 2010.
129. R. Fergus, L. Fei-Fei, P. Perona and A. Zisserman. Learning object categories from Internet image searches. *Proc. Of IEEE, Special Issue on Internet Vision. Vol 98/8*. 2010.
130. B. Yao and L. Fei-Fei. Modeling mutual context of object and human pose in human-object interaction activities. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2010.
131. B. Yao and L. Fei-Fei. Grouplet: a structured image representation for recognizing human and object interactions. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2010.
132. L.J. Li, C. Wang, Y. Lim, D. Blei and L. Fei-Fei. Building and Using a Semantivisual image hierarchy. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2010.
133. J.C. Niebles, L. Chen and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2010.
134. R. Socher and L. Fei-Fei. Connecting Modalities: Semi-supervised Segmentation and Annotation of Images Using Unaligned Text Corpora. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2010.
135. D. B. Walther, D. M. Beck, and L. Fei-Fei. To err is human: investigating neural function by correlating error patterns with human behavior. in: Nikolaus Kriegeskorte and Gabriel Kreiman (eds.), *Understanding visual population codes – Toward a common multivariate framework for cell recording and functional imaging*, MIT Press, Cambridge, Massachusetts. 2010.

136. B. Yao, D.B. Walther, D.M. Beck*, L. Fei-Fei*. Hierarchical Mixture of Classification Experts Uncovers Interactions between Brain Regions. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2009. (* indicates equal contribution)
137. B. Chai†, D.B. Walther†, D.M. Beck*, L. Fei-Fei*. Exploring Functional Connectivity of the Human Brain using Multivariate Information Analysis. *Proceedings of the Neural Information Processing Systems (NIPS)*. 2009. (†,* indicates equal contribution)
138. L.-J. Li and L. Fei-Fei. OPTIMOL: automatic Online Picture collection via Incremental Model Learning. *International Journal of Computer Vision (IJCV)*. 2009.
139. H. Su, M. Sun, S. Savarese and L. Fei-Fei. A Multi-View Probabilistic Model for 3D Object Classes. *IEEE Inter. Conf. Comp. Vision (ICCV)*. 2009
140. M. Peelen, L. Fei-Fei, and S. Kastner. Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature*. doi:10.1038/nature08103. 2009.
141. D. B. Walther, E. Caddigan, L. Fei-Fei*, and D. M. Beck*. Natural scene categories revealed in distributed patterns of activity in the human brain. *Journal of Neuroscience*, 29(34):10573-10581, 2009.
142. L.-J. Li, R. Socher and L. Fei-Fei. Towards Total Scene Understanding: Classification, Annotation and Segmentation in an Unsupervised Framework. *IEEE Inter. Conf. Comp. Vision and Pattern Recog (CVPR)*. 2009
143. C. Wang, D. Blei and L. Fei-Fei. Simultaneous Image Classification and Annotation. *IEEE Inter. Conf. Comp. Vision and Pattern Recog*. 2009
144. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei. ImageNet: A Preview of a Large-Scale Hierarchical Database. *IEEE Inter. Conf. Comp. Vision and Pattern Recog (CVPR)*. 2009
145. M. Sun, H. Su, S. Savarese and L. Fei-Fei. A Multi-View Probabilistic Model for 3D Object Classes. *IEEE Inter. Conf. Comp. Vision and Pattern Recog (CVPR)*. 2009
146. J.C. Niebles, B. Han, A. Ferencz and L. Fei-Fei. Extracting Moving Humans from Internet Videos. *IEEE European Conf. Computer Vision (ECCV)*, 2008.
147. B. Collins, J. Deng, K. Li and L. Fei-Fei. Towards scalable dataset construction: An active learning approach. *IEEE European Conf. Computer Vision (ECCV)*, 2008.
148. S. Savarese and L. Fei-Fei. View synthesis for recognizing unseen poses of object classes. *IEEE European Conf. Computer Vision (ECCV)*, 2008.
149. J.C. Niebles, H. Wang and L. Fei-Fei, Unsupervised Learning of Human Action Categories Using Spatial-Temporal Words. *International Journal of Computer Vision (IJCV)*, DOI 10.1007/s11263-007-0122-4, 2008.
150. S. Savarese, A. Del Pozo, J.C. Niebles, L. Fei-Fei, "Spatial-Temporal Correlators for Unsupervised Action Classification", *IEEE Workshop on Motion and Video Computing*, Copper Mountain, Colorado January 8-9, 2008.
151. S. Savarese and L. Fei-Fei. 3D generic object categorization, localization and pose estimation. *IEEE Intern. Conf. in Computer Vision (ICCV)*. 2007.
152. L. J. Li and L. Fei-Fei. What, where and who? Classifying event by scene and object recognition. *IEEE Intern. Conf. in Computer Vision (ICCV)*. 2007.
153. L. Cao and L. Fei-Fei. Spatially coherent latent topic model for concurrent object segmentation and classification. *IEEE Intern. Conf. in Computer Vision (ICCV)*. 2007.
154. D. B. Walther and L. Fei-Fei. Measuring the cost of deploying top-down visual attention. *Journal of Vision*, 7(11):9, 1-12, <http://journalofvision.org/7/11/9/>, doi:10.1167/7.11.9. 2007.
155. L. Fei-Fei, R. Fergus and P. Perona. Learning generative visual models for 101 object categories. *Computer Vision and Image Understanding (CVIU)*. 2007.

156. L.J. Li, J.C. Niebles and L. Fei-Fei. OPTIMOL: a framework for Online Picture collecTion via Incremental MOdel Learning. *Association for the Advancement of Artificial Intelligence (AAAI) 2007 Robot Competition and Exhibition*, 2007
157. L.J. Li, G. Wang and L. Fei-Fei. OPTIMOL: automatic Object Picture collecTion via Incremental MOdel Learning. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2007. ((*))
158. J.C. Niebles and L. Fei-Fei. A hierarchical model of shape and appearance for human action classification. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2007.
159. L. Fei-Fei, Iyer, A., Koch, C., & Perona, P. What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1):10, 1-29, <http://journalofvision.org/7/1/10/>, doi:10.1167/7.1.10. 2007.
160. J.C. Niebles, H. Wang, L. Fei-Fei. Unsupervised learning of human action categories using spatial-temporal words. *British Machine Vision Conference (BMVC) 2006*.
161. A. Kushal, M. Rahrkar, J. Ponce, T. Huang and L. Fei-Fei. Audio-Visual Speaker Localization Using Graphical Models. *Int. Conf. Patt. Recog.* 2006.
162. J. Tu, A. Ivanovic, T. Huang, and L. Fei-Fei. Variational Shift Invariant Probabilistic PCA for Face Recognition. *Int. Conf. Patt. Recog.* 2006
163. L. Fei-Fei. Knowledge transfer in learning to recognize visual object classes. *IEEE ICDL*. 2006
164. G. Wang, Y. Zhang, and L. Fei-Fei. Using dependent local regions for categorizing objects in a generative framework. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2006.
165. L. Fei-Fei, R. Fergus and P. Perona. One-Shot Learning of 100 Object Categories. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*. Vol28(4), 594 - 611, 2006.
166. R. VanRullen, L. Reddy and L. Fei-Fei. Binding for natural objects. *Vision Research*. Vol45(25-26), 3133-3144. 2005.
167. L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models for 101 object categories. *Computer Vision and Image Understanding (CVIU)*, 2006.
168. R. Fergus, L. Fei-Fei, P. Perona and A. Zisserman. Learning Object Categories from Google's Image Search. *IEEE Intern. Conf. in Computer Vision (ICCV)*. 2005.
169. L. Fei-Fei and P. Perona. A Bayesian Hierarchical Model for Learning Natural Scene Categories. *IEEE Computer Vision and Pattern Recognition (CVPR)*. 2005.
170. L. Fei-Fei, R. VanRuellen, C. Koch and P. Perona. Why does natural scene categorization require little attention? Exploring attentional requirements for natural and synthetic stimuli. *Visual Cognition*. 2005.
171. L. Fei-Fei, R. VanRullen, C. Koch and P. Perona. Why does natural scene recognition require little attention? Exploring attentional requirements for natural and synthetic stimuli. *Visual Cognition*. 12(6): pp893-924. 2005.
172. L. Fei-Fei, R. Fergus and P. Perona. Learning generative visual models from few training examples: an incremental Bayesian approach tested on 101 object categories. *IEEE Computer Vision and Pattern Recognition (CVPR), Workshop on Generative-Model Based Vision*. 2004.
173. S. Savarese, L. Fei-Fei and P. Perona, "What do reflections tell us about the shape of a mirror?" in *Applied Perception in Graphics and Visualization [sponsored by ACM SIGGRAPH], Los Angeles, August 7-8, 2004*.
174. L. Fei-Fei, R. Fergus and P. Perona. A Bayesian approach to unsupervised One-Shot learning of Object categories. *Proc. International Conference on Computer Vision (ICCV)*. 2003.
175. F.F. Li, R. VanRullen, C. Koch and P. Perona. Rapid natural scene categorization in the near absence of attention. *Proc. Natl. Acad. Sci.* 99, 8378 – 8383, 2002. ((*))
176. G.B. Stanley, F.F. Li, and Y. Dan. Reconstruction of natural scenes from ensemble responses in the LGN, *Journal of Neuroscience*, 19(18):8036-8042, 1999.

PATENTS

1. Fei-Fei Li, Jia Deng, Jonathan Krause & Alexander C. Berg. Method and System for Optimizing Accuracy-Specificity Trade-offs in Large Scale Visual Recognition, Stanford Docket #S12-199, patent pending.
2. Olga Russakovsky, Yuanqing Lin, Kai Yu, & Fei-Fei Li. Object-Centric Spatial Pooling for Image Classification. US 8,761,510 B2

PROFESSIONAL AFFILIATIONS

2006 – present Senior Member, IEEE Computer Society (IEEE)

2013 – present Member, Association of Computing Machinery (ACM)

CONFERENCE ORGANIZING COMMITTEE

WORKSHOP, TUTORIAL AND COURSE ORGANIZER OR CO-CHAIR

CONFERENCE ORGANIZING COMMITTEE

- General Chair, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2016
- Local Arrangement Chair, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2010
- Industrial Liaison, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2012
- Area Chair, Neural Information Processing Systems (NIPS), 2010, 2011, 2013
- Area Chair, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2009, 2010, 2012, 2013
- Area Chair, IEEE International Conference of Computer Vision (ICCV), 2007, 2011
- Area Chair, European Conference of Computer Vision (ECCV), 2012

WORKSHOP, TUTORIAL AND COURSE ORGANIZER OR (CO-)CHAIR

- Co-chair, Stanford-White House Office of Science and Technology Policy (OSTP) Symposium on The Future of Artificial Intelligence, 2016
- Co-organizer, ImageNet Large-Scale Visual Recognition Challenge, 2010 (ECCV), 2011 (ICCV), 2012 (ECCV), 2013 (ICCV), 2014 (ECCV)
- Co-organizer, Stanford Workshop on AI and Knowledge (SWANK), 2014
- Co-organizer, Fine-Grained Categorization Challenge, 2013
- Co-organizer, Workshop on “Big Data Meets Computer Vision”, Neural Information Processing Systems (NIPS), 2012
- Co-chair, Mechanical Turk for Computer Vision, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2010
- Organizer, Bay Area Vision Meeting (BAVM), 2009, 2012
- Co-director, 1st Sino-USA Summer School in Vision Learning and Pattern Recognition (VLPR), 2009
- Co-chair, 1st International Workshop in Visual Scene Understanding (ViSU), IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2009
- Co-chair, International Workshop in 3D Representation for Recognition, IEEE International Conference of Computer Vision (ICCV), 2007
- Co-organizer, Tutorial on Recognizing and Learning Object Categories, IEEE Conference of Computer Vision and Pattern Recognition (CVPR), 2007
- Co-chair, Vision Sciences Society (VSS) Annual Meeting Satellite Symposium: Natural Scene Understanding: Statistics, Recognition and Representation, 2007
- Co-chair, International Workshop in Computer Vision (in collaboration with Microsoft MSRA), Tibet, 2006

- Co-organizer, Tutorial on Recognizing and Learning Object Categories, IEEE International Conference on Computer Vision (ICCV), 2005

EDITORIAL BOARD AND JOURNAL REVIEWER

- Associate Editor: IEEE Transaction in Pattern Analysis and Machine Intelligence (PAMI)
- Reviewer:
 - International Journal of Computer Vision (IJCV)
 - IEEE Transaction in Pattern Analysis and Machine Intelligence (PAMI)
 - Computer Vision and Image Understanding (CVIU)
 - ACM Transaction on Applied Perception (TAP)
 - Journal of Machine Learning Research (JMLR)
 - Nature
 - Cognitive Psychology
 - Journal of Vision (JoV)
 - NeuroImage
 - Journal of Experimental Psychology (JEP)
 - ACM Transaction on Applied Perception (TAP)
 - Cerebral Cortex

ADVISORY BOARD MEMBER OF ACADEMIC PROGRAMS OR INSTITUTIONS

- Scientific Advisory Board (Fachbeirat), Max Planck Institute of Informatics in Saarbrucken, Germany, 2017 - 2020
- External Advisory Committee, Center for Brains, Minds and Machines, MIT, 2017 - 2019

TEACHING

STANFORD UNIVERSITY

CS231n: Convolutional Neural Network for Visual Recognition

- Spring 2017, Enrollment: 770
- Winter 2016, Enrollment: 320
- Winter 2015, Enrollment: 156

CS131: Computer Vision: Foundations and Applications

- Fall 2016, Enrollment: 100
- Fall 2015, Enrollment: 80
- Fall 2014, Enrollment: 49
- Fall 2013, Enrollment: 35

CS231a (previous numbered CS223b): Computer Vision

- Fall 2012, Enrollment: 80
- Fall 2011, Enrollment: 110
- Winter 2010, Enrollment: 130

CS331 (previous numbered CS323): Advanced Reading in Computer Vision

- Winter 2014, Enrollment: 8
- Spring 2013, Enrollment: 18
- Fall 2009, Enrollment: 12

CS431(previous numbered CS423b)/PSY250: High-Level Scene Understanding

- Spring 2014, co-instructor Prof. Kalanit Grill-Spector, Enrollment: 19
- Spring 2010, co-instructor Prof. Kalanit Grill-Spector, Enrollment: 20

PRINCETON UNIVERSITY

COS429: Computer Vision

- Fall 2008, Enrollment: 30
- Fall 2007, Enrollment: 28

COS598: High-Level Recognition in Computer Vision

- Spring 2007, Enrollment: 14

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

ECE546: Computer Vision

- Fall 2006, Enrollment: 45

ECE598: Readings in Computer Vision and Learning

- Fall 2005, Enrollment: 25

STUDENTS / ADVISING

POSTDOCTORAL SCHOLARS

2016. 08 – 2018.08	Animesh Garg (Ph.D. from U. of California Berkeley) Placement: Research Scientist, NVidia; Assistant Professor, U. Toronto (deferred)
2016. 08 – 2017.05	Judy Hoffman (Ph.D. from U. of California Berkeley) Placement: Postdoc, U.C. Berkeley
2015. 07 – 2016.12	Joseph Lim (Ph.D. from MIT) Placement: Assistant Professor, University of Southern California
2014.10 – 2016	Lamberto Ballan (Ph.D. from U. of Florence, Italy) Placement: Research scientist, University of Florence, Italy
2013.01 – 2017.08	Guido Pusiol (Ph.D. from INRIA) Placement: Startup
2011.06 – 2017.06	Michelle Greene (Ph.D. from MIT) Placement: Assistant Professor, Bain College
2012.09 – 2017.08	Alexandre Alahi (Ph.D. from EPFL) Placement: Assistant Professor, EPFL
2013.01 – 2015	Armand Joulin (Ph.D. from Ecole Polytechnique) Placement: Research Scientist, Facebook Research Lab
2013.07 – 2014.09	Alireza Fathi (Ph.D. from Georgia Tech) Placement: Research Scientist, Apple, Inc.
2012.06 – 2013.10	Jia Deng (Ph.D. from Princeton) Placement: Assistant Professor, U. Michigan Ann Arbor
2006 – 2010	Dirk B. Walther (Ph.D. from California Institute of Technology), UIUC Placement: Assistant Professor, U. of Toronto, Canada

PH.D. STUDENTS

- 2016 -- present Jim Fan (CS, Stanford)
- 2016 -- present Danfei Xu (CS, Stanford; co-advised by Silvio Savarese)
- 2016 – present Albert Haque (CS, Stanford)
- 2016 – present Ranjay Krishna (CS, Stanford)
- 2015 – present De-An Huang (CS, Stanford)
- 2015 – present Yuke Zhu (CS, Stanford)
- 2013 – 2018 Justin Johnson (CS, Stanford)
Dissertation: “*Compositional Visual Intelligence*”
Placement: Assistant Professor, U. Michigan; Research Scientist, Facebook
- 2013 – 2018 Serena Yeung (EE, Stanford)
Dissertation: “*Visual Understanding of Human Activity: Towards Ambient Intelligence in AI-assisted Hospitals*”
Placement: Assistant Professor, Stanford University
- 2013 – 2017 Timnit Gebru (EE, Stanford)
Dissertation: “*Visual Computational Sociology: Computer Vision Methods and Challenges*”
Placement: Postdoc, Microsoft Research NY
- 2012 – 2016 Andrej Karpathy (CS, Stanford)
Dissertation: “*Connecting images with natural language*”
Placement: Founding Research Scientist at OpenAI
- 2012 – 2016 Vignesh Ramanathan (EE, Stanford)
Dissertation: “*Human-centric video understanding with weak supervision*”
Placement: Research Scientist at Facebook AI Research Lab (FAIR)
- 2011 – 2016 Jon Krause (CS, Stanford)
Dissertation: “*Fine-grained object recognition*”
Placement: Research Scientist at Google Brain
- 2010 – 2015 Kevin Tang (CS, Stanford)
Dissertation: “*Visual learning with weakly labeled video.*”
Placement: Research Scientist at Snapchat
- 2009 – 2016 Marius Catalin Iordan (CS, Stanford)
Dissertation: “*Uncovering the Neural Representation of Multiple Dimensions of Object Categorization in Human Visual Cortex.*”
Placement: Postdoctoral Scholar at Princeton University
- 2009 – 2015 Olga Russakovsky (CS, Stanford)
Dissertation: “*Scaling Up Object Detection*”
Placement: Postdoctoral Scholar at Carnegie Mellon University
- 2009 – 2015 Chris Baldassano (CS, Stanford)
Dissertation: “*Visual scene perception in the human brain: Connections to memory, categorization, and social cognition*”
Placement: Postdoctoral Scholar at Princeton University
- 2008 – 2013 Bangpeng Yao (CS, Stanford)
Dissertation: “*Understanding Human Actions in Still Images*”
Placement: Algorithm Developer at Hudson River Trading
- 2007 – 2012 Jia Deng (CS, Princeton)

- 2006 – 2011
 Dissertation: “*Large Scale Visual Recognition*”
 Placement: Assistant Professor at U. Michigan, Ann Arbor
 Jia Li (CS, Stanford)
- 2005 – 2010
 Dissertation: “*Semantic Image Understanding from the Web, in Large Scale and with Real-World Data*”
 Placement: Research Scientist at Yahoo! Research
 Juan Carlos Niebles (EE, Princeton)
- Dissertation: “*Extracting Moving People and Categorizing Their Activities in Videos*”
 Placement: Assistant Professor at Universidad del Norte, Colombia

MASTER STUDENTS

- 2013 – 2016 Yuke Zhu (CS, Stanford)
 Placement: Ph.D. student at Stanford University
- 2014 – 2016 Ranjay Krishna (CS, Stanford)
 Placement: Ph.D. student at Stanford University
- 2014 – 2015 Sean Ma (CS, Stanford)
- 2010 – 2011 Aditya Khosla (CS, Stanford)
 Placement: Ph.D. student at MIT
- 2011 – 2013 Yongwhan Lim (CS, Stanford)
 Placement: Ph.D. student at MIT
- 2010 – 2011 Vinayak Agarwal (CS, Stanford)
- 2010 – 2011 Rob Cosgrill (EE, Stanford)
 Placement: VideoSurf, Inc.
- 2009 – 2010 Georgios Georgiadis (CS, Stanford)
 Placement: Ph.D. student at UCLA
- 2010 – 2011 Dan Goodwin (EE, Stanford)
 Placement: IDEO, Inc
- 2009 – 2011 Yu Lou (CS, Stanford)
 Placement: A9.com
- 2010 – 2012 Sanjeev Satheesh (CS, Stanford)
 Placement: Microsoft
- 2011 – 2012 Jiahu Shi (CS, Stanford)
- 2010 – 2012 Ningxuan Wang (CS, Stanford)
 Placement: Microsoft
- 2010 – 2012 Haizi Yu (Statistics, Stanford)

UNDERGRADUATE STUDENTS

- 2014 – 2016 Emily Tang (CS, Stanford)
- 2014 – 2016 Clara Fanning-Jiang (CS, Stanford)
- 2009 – 2011 Yongwhan Lim (CS, Stanford)
- 2007 – 2008 Bredan Collins (CS, Princeton)
- 2009 Minh Do (EE, Princeton)
- 2007 – 2008 Paul Kompfer (CS, Princeton)
- 2007 – 2008 Douglas Horhensee (CS, Princeton)

2006 – 2007

Brian Leo Quanz (ECE, UIUC)

2006 – 2007

Justos Birgiolos (ECE, UIUC)

ADDITIONAL INFORMATION

Languages: Mandarin, English