

LIANGLIANG CAO

🌐 www.llcao.net ✉ liangliang.cao@gmail.com

01/31/2022

EMPLOYMENT

Senior Staff Research Scientist and Senior Manager, Google 10/2018 - now

EDUCATION

Ph.D., University of Illinois at Urbana-Champaign 2011

M.Phil., Multimedia Lab, The Chinese University of Hong Kong 2005

B. E., Dept. Electronic Engineering, University of Science and Technology of China 2003

WORK HISTORY

Affiliated Research Associated Professor, UMass Amherst 08/2018 - 06/2021

Co-founder and CTO, Switi Inc (acquired) 08/2016 - 10/2018

Adjunct Associate Professor, Columbia University in the City of New York 01/2013 - 06/2018

Senior Research Scientist, Yahoo Labs at New York City 04/2015 - 08/2016

Research Staff Member, IBM T. J. Watson Research Center, NY 06/2011 - 04/2015

RESEARCH INTERESTS

Computer Vision, Speech and Language, Multimedia, Cloud AI

EXPERIENCES AND PROJECTS

Google Cloud Vision 11/2021-now

Senior Staff Research Scientist

- Tech Lead for Cloud vision modeling quality

Google Speech 02/2020-11/2021

Senior Staff Research Scientist and Tech Lead Manager

- Built a team of 10 engineers and scientists from scratch (Google people manager score consistently higher than 95%).
- Developed the cutting-edge End2end Speech Recognition models for Google Cloud.
- Reduced the word error rates of Cloud speech recognition by 50%.
- Deployed speech models for Google Cloud APIs.
- Launched Google's On-Premise Speech2Text, Google's first AI product on Anthos.

Google Contact Center AI 10/2018-02/2020

Staff Research Scientist

- Acted as the tech lead to win the Verizon speech competition.
- Won the largest contract in the history of Google Cloud AI.
- Promoted to Senior Staff Research Scientist.

Switi Inc. 08/2016-10/2018

Co-founder and CTO

- Built AI assistants for enterprise customers.
- Worked as architects, product manager, and sales engineer.
- Helped the company to be acquired in 2018.

Yahoo Inc. 04/2016-08/2018

Senior Research Scientist

- Helped to build the computer vision team in Yahoo Lab NYC.

- Co-led the launch of Yahoo Taiwan's mobile image search

IBM Watson Research Center

06/2011-03/2015

Research Staff Member

- Proposed and implement the idea of making IBM Multimedia Analytics and Search (IMARS) 100x faster

SELECTED AWARDS

- Google Cloud Feats of Engineering Award, 2019
- ACM SIGMM Rising Star Award, 2017
- IBM Outstanding Accomplishment for Multimedia Team, 2012
- Best Paper Award, SIGKDD Workshop on BigMine, 2012
- IBM Research Division Awards, 2011 - 2013
- IBM Watson Emerging Leader in Multimedia and Signal Processing, 2010
- First Place in ImageNet Large Scale Visual Recognition Challenge, 2010

PUBLICATION SUMMARY

Google Scholar: <http://scholar.google.com/citations?user=S-hBSfIAAAAJ&hl=en>

Total 7200+ citations, H-Index 42.

I mainly published in Computer Vision (including 11 CVPR, 5 ICCV, 3 TPAMI and 1 PIEEE) and Multimedia (including 8 ACM-MM, 3 WWW and 5 TMM). During 2020-2021, I was in charge of Cloud speech modeling and focused on the Speech field (published 4 ICASSP and 4 Interspeech papers).

SELECTED PUBLICATIONS

Journals

17. Y. Zhang, D. Park, W. Han, J. Qin, A. Gulati, J. Shor, A. Jansen, Y. Xu, Y. Huang, S. Wang, Z. Zhou, B. Li, M. Ma, W. Chan, J. Yu, Y. Wang, L. Cao, K. C. Sim, B. Ramabhadran, T. Sainath, F. Beaufays, Z. Chen, Q. V. Le, C.-C. Chiu, R. Pang, Y. Wu, BigSSL: Exploring the Frontier of Large-Scale Semi-Supervised Learning for Automatic Speech Recognition, *in submission*, 2021.
16. Y. Liang, L. Jiang, L. Cao, Y. Kalantidis, L.-J. Li, A. Hauptmann, Focal Visual-Text Attention for Memex Question Answering, *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, Vol.41, No. 8, pp. 1893-1908, 2019.
15. Y. Li, L. Cao, J. Zhu, J. Luo, Mining Fashion Outfit Composition Using An End-to-End Deep Learning Approach on Set Data, *IEEE Transaction on Multimedia*, Vol. 9, No. 8, pp. 1946-1955, 2017.
14. L. Wang, X. Zhao, Y. Si, L. Cao, Y. Liu. Context-associative Hierarchical Memory Model for Human Activity Recognition and Prediction. *IEEE Transaction on Multimedia*, Vol. 19, No. 3: pp. 646-659, 2017.
13. Q. You, R. Pan, L. Cao, J. Luo, Image Based Appraisal of Real Estate Properties, *IEEE Transaction on Multimedia*, Vol 19, No. 8, pp. 1946-1955, 2017.
12. Q. You, L. Cao, Y. Cong, X. Zhang, and J. Luo, A Multifaceted Approach to Social Multimedia-based Prediction of Elections, *IEEE Transaction on Multimedia*, Vol 17, No. 12, 2015.
11. J. Smith, L. Cao, N. Codella, M. Hill, M. Merler, Q. Nguyen, E. Pring, R.A. Uceda-Sosa, Massive-scale Learning of Image and Video Semantic Concepts, *IBM Journal of Research and Development*, Vol. 59, Issue 2/3, 2015.
10. W. Guo, L. Cao, T. Han, S. Yan, and C. Xu, Max-Confidence Boosting with Uncertainty for Visual Tracking, *IEEE Transaction on Image Processing*, Vol. 24, No. 5, pp. 1650-1659, 2015.
9. G.-J. Qi, S.-F. Tsai, M.-H. Tsai, L. Cao, and T. S. Huang, Web-Scale Multimedia Information Networks. *Proceedings of the IEEE (PIEEE)*, Vol. 100, No. 9, p. 2688-2704, 2012.
8. Z. Yin, L. Cao, Q. Gu, J. Han, Latent Community Topic Analysis: Integration of Community Discovery with Topic Modeling, *ACM Transactions on Intelligent Systems and Technology*, vol. 3, No. 4, 2012.
7. Y. Tian, L. Cao, Z. Liu, and Z. Zhang: Hierarchical Filtered Motion for Action Recognition in Crowded Videos. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 42, No.3: 313-323, 2012.

6. L. Cao, X. Jin, Z. Yin, A. Del Pozo, J. Luo, J. Han and T. S. Huang, RankCompete: Simultaneous Ranking and Clustering of Information Networks, *Neurocomputing*, vol. 95, p. 98-104, 2011.
5. X. Jin, S. Kim, J. Han, L. Cao, and Z. Yin: A General Framework for Efficient Clustering of Large Datasets Based on Activity Detection. *Statistical Analysis and Data Mining*, vol. 4, No. 1: 11-29, 2011.
4. S. Chen, L. Cao, Y. Wang, J. Liu, and X. Tang, Image Segmentation by MAP-ML Estimations, *IEEE Transactions on Image Processing*, Vol 19, No. 9, p. 2254-2264, 2010.
3. L. Cao, J. Luo, H. Kautz and T. S. Huang, Image Annotation within the Contest of Personal Photo Collections, *IEEE Transactions on Multimedia*, Vol. 11, No. 2, p. 208-219, 2009.
2. J. Liu, L. Cao, Z. Li and X. Tang. Plane-Based Optimization for 3D Object Reconstruction from Single Line Drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 30 (2), p. 315-327, 2008.
1. L. Cao, J. Liu and X. Tang. What the Back of the Object Looks Like: 3D Reconstruction from Line Drawings Without Hidden Line. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 30 (3), p. 507-517, 2008.

Conferences

83. Z. Lu, Y. Pan, T. Doutre, L. Cao, R. Prabhavalkar, C. Zhang, T. Strohman, Input Length Matters: An Empirical Study Of RNN-T And MWER Training For Long-form Telephony Speech Recognition, *in submission*, 2022
82. Q. Li, Y. Zhang, D. Qiu, Y. He, L. Cao, P. C. Woodland, Improving Confidence Estimation on Out-of-Domain Data for End-to-End Speech Recognition, *ICASSP*, 2022
81. T. Doutre, W. Han, C.-C. Chiu, R. Pang, O. Siohan, L. Cao: Bridging the gap between streaming and non-streaming ASR systems by distilling ensembles of CTC and RNN-T models, *Interspeech*, 2021.
80. Q. Li, Y. Zhang, B. Li, L. Cao, P. C. Woodland: Residual Energy-Based Models for End-to-End Speech Recognition, *Interspeech*, 2021.
79. Z. Lu, W. Han, Y. Zhang, L. Cao, Exploring Targeted Universal Adversarial Perturbations to End-to-end ASR Models, *Interspeech*, 2021.
78. D. Qiu, Y. He, Q. Li, Y. Zhang, L. Cao, I. McGraw, Multi-Task Learning for End-to-End ASR Word and Utterance Confidence with Deletion Prediction, *Interspeech*, 2021.
77. T. Doutre, W. Han, M. Ma, Z. Lu, C.-C. Chiu, R. Pang, A. Narayanan, A. Misra, Y. Zhang, L. Cao: Improving Streaming Automatic Speech Recognition With Non-Streaming Model Distillation On Unsupervised Data, *ICASSP*, 2021.
76. Q. Li, D. Qiu, Y. Zhang, B. Li, Y. He, P. C. Woodland, L. Cao, T. Strohman: Confidence Estimation for Attention-based Sequence-to-sequence Models for Speech Recognition, *ICASSP*, 2021.
75. D. Qiu, Q. Li, Y. He, Y. Zhang, B. Li, L. Cao, R. Prabhavalkar, D. Bhatia, W. Li, K. Hu, T. Sainath, I. McGraw: Learning Word-Level Confidence For Subword End-to-End ASR, *ICASSP*, 2021.
74. C.-C. Chiu, A. Narayanan, W. Han, R. Prabhavalkar, Y. Zhang, N. Jaitly, R. Pang, T. Sainath, P. Nguyen, L. Cao, Y. Wu: RNN-T Models Fail to Generalize to Out-of-Domain Audio: Causes and Solutions, *IEEE SLT*, 2020.
73. Z. Lu, L. Cao, Y. Zhang, C.-C. Chiu, J. Fan, Speech Sentiment Analysis via Pre-Trained Features from End-to-End ASR Models, *ICASSP 2020*: 7149-7153.
72. M. Gadelha, A. RoyChowdhury, G. Sharma, E. Kalogerakis, L. Cao, E. Learned-Miller, R. Wang, S. Maji: Label-Efficient Learning on Point Clouds using Approximate Convex Decompositions, *ECCV*, 2020.
71. J. Liu, L. Cao, Y. Tian: Active Learning for Effective Pulmonary Nodule Detection. *MICCAI 2020*.
70. J. Liu, L. Cao, O. Akin, Y. Tian: 3DFPN-HS2: 3D Feature Pyramid Network Based High Sensitivity and Specificity Pulmonary Nodule Detection. *MICCAI 2019*.

69. A. RoyChowdhury, P. Chakrabarty, A. Singh, S.-Y. Jin, H. Jiang, L. Cao, E. Learned-Miller: Automatic Adaptation of Object Detectors to New Domains Using Self-Training. *CVPR* 2019: 780-790
68. Z. Shen, H. Shi, J. Yu, H. Phan, R. Feris, L. Cao, D. Liu, X. Wang, T. S. Huang, M. Savvides: Improving Object Detection from Scratch via Gated Feature Reuse. *BMVC* 2019.
67. J. Liang, L. Jiang, L. Cao, A. Hauptmann, Answer with Grouping Snippets: Focal Visual-Text Attention for Visual Question Answering, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. (Spotlight)
66. Y. Lu, H. Lu, L. Cao, F. Wu, Learning Deterministic Policy with Target for Inter-Cell Interference Coordination in Wireless Networks, *IEEE Global Communication Conference (GlobeCom)*, 2018
65. W. Tan, S. Chang, L. Fong, C. Li, Z. Wang, L. Cao, Matrix Factorization on GPUs with Memory Optimization and Approximate Computing, *International Conference on Parallel Processing (ICPP)*, 2018.
64. H. Akbari, H. Arora, L. Cao, N. Mesgarani, Lip2Audspec: Speech Reconstruction from Silent Lip Movements Video, *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
63. Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, J. Li, Learning from Noisy Labels with Distillation, *IEEE International Conference on Computer Vision (ICCV)*, 2017.
62. L. Jiang, Y. Kalantidis, L. Cao, S. Farfadi, J. Tang and A. Hauptmann, Delving Deep into Personal Photo and Video Search, *International Conference on Web Search and Data Mining (WSDM)*, 2017.
61. L. Jiang, L. Cao, Y. Kalantidis, S. Farfadi, A. Hauptmann, Visual Memory Question Answering, *AAAI*, 2017. (demo)
60. W. Tan, L. Cao, L. Fong, Faster and Cheaper: Parallelizing Large-Scale Matrix Factorization on GPUs, *ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2016. (also appeared in NVIDIA Research Highlight)
59. W. Ma, L. Cao, L. Yu, G. Long and Y. Li, GPU-FV: Realtime Fisher Vector and Its Applications in Video Monitoring, *ACM International Conference on Multimedia Retrieval (ICMR)*, 2016.
58. L. Cao, J. Hsiao, P. de Juan, Y. Li and B. Thomee, Incremental Learning for Fine-Grained Image Recognition, *ACM International Conference on Multimedia Retrieval (ICMR)*, 2016.
57. C. Wang, L. Cao, J. Fan, Building Joint Spaces for Relation Extraction, *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
56. Y. Bai, W. Ma, Y. Li, L. Cao, W. Guo, L. Yang, Multi-Scale Fully Convolutional Network for Fast Face Detection, *British Machine Vision Conference (BMVC)*, 2016
55. Q. You, L. Cao, H. Jin, J. Luo, Robust Visual-Textual Sentiment Analysis: When Attention meets Tree-structured Recursive Neural Networks, *ACM Conference on Multimedia (ACM MM)*, 2016.
54. R. Schifanella, P. de Juan, J. Tetreault, L. Cao, Detecting Sarcasm in Multimodal Social Platforms, *ACM Conference on Multimedia (ACM MM)*, 2016.
53. Y. Li, Y. Song, L. Cao, J. Tetreault, J. Luo, TGIF: A New Dataset and Benchmark on Animated GIF Description, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. (Spotlight)
52. M. Gygli, Y. Song, L. Cao, Video2GIF: Automatic Generation of Animated GIFs from Video, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
51. N. Yakovenko, L. Cao, C. Raffel, and J. Fan, Poker-CNN: A Pattern Learning Strategy for Making Draws and Bets in Poker Games, *AAAI*, 2016.
50. C. Wang, L. Cao, B. Zhou, Medical Synonym Extraction with Concept Space Models, *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2015.
49. J. Bowler, R. Feris, L. Cao, J. Wang, M. Zhou, Automated Axon Segmentation from Highly Noisy Microscopic Videos, *IEEE Winter Conference on Applications of Computer Vision (WACV)* 2015.

48. M. Merler, L. Cao, J. R. Smith, You are what you tweet... pic! gender prediction based on semantic analysis of social media images, *IEEE International Conference on Multimedia and Expo (ICME)*, 2015.
47. L. Cao, Z. Yan, J. R. Smith, Multi-facet Learning using Deep Convolutional Neural Network for Person-Related Categories in Photos, *ACM International Conference on Multimedia Retrieval (ICMR)*, 2015.
46. Y. Cui, Y. Xiang, K. Rong, R. Feris, L. Cao, A Spatial-Color Layout Feature for Content-based Galaxy Image Retrieval, *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014.
45. S. Albaradei, Y. Wang, L. Cao, L.-J. Li: Learning mid-level features from object hierarchy for image classification, *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014.
44. Y. Bao, L. Cao, H. Li, J. Tang, Cuteness Recognition and Localization in the Photos of Animals, *ACM Conference on Multimedia (ACM MM)*, 2014.
43. F. Yu, L. Cao, M. Merler, N. Codella, T. Chen, J. R. Smith; S.-F. Chang, Modeling attributes from category-attribute proportions, *ACM Conference on Multimedia (ACM MM)*, 2014.
42. Q. Chen, Z. Song, R. Feris, A. Datta, L. Cao, Z. Huang, S. Yan, Efficient Maximum Appearance Search for Large-Scale Object Detection, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
41. F. Yu, L. Cao, R. Feris, J. R. Smith, S.-F. Chang, Designing Category-Level Attributes for Discriminative Visual Recognition, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
40. Z. Li, S. Chang, F. Liang, T. S. Huang, L. Cao, J. R. Smith, Learning Locally-Adaptive Decision Functions for Person Verification *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
39. X. Liang, L. Lin, and L. Cao, Learning Latent Spatio-Temporal Compositional Model for Human Action Recognition, *ACM Conference on Multimedia (ACM MM)*, 2013. (Oral presentation)
38. L. Cao, L. Gong, J. R. Kender, N. Codella, J. R. Smith, Learning by Focusing: A New Framework for Concept Recognition and Feature Selection, *IEEE International Conference on Multimedia and Expo*, 2013.
37. N. Codella, G. Hua, L. Cao, M. Merler, L. Gong, M. Hill, J. Smith: Large-scale video event classification using dynamic temporal pyramid matching of visual semantics, *IEEE International Conference on Image Processing (ICIP)*, 2013.
36. L. Cao, Z. Li, Y. Mu, S.-F. Chang, A Unified Submodular Framework towards Video Pooling and Hashing, *ACM Conference on Multimedia (ACM MM)*, 2012. (Oral presentation)
35. L. Cao, Y. Mu, A. Natsev, S.-F. Chang, G. Hua, J. R. Smith, Scene Aligned Pooling for Complex Video Recognition, *European Conference on Computer Vision (ECCV)*, 2012.
34. L. Cao, J. Smith, Z. Wen, Z. Yin, X. Jin, J. Han, BlueFinder: Estimate Where a Beach Photo Was Taken, *International Conference on World Wide Web (WWW)*, 2012.
33. L. Cao, H. D. Kim, M.-H. Tsai, B. Cho, Z. Li, I. Gupta, C. Zhai, and T. S. Huang, Delta-SimRank Computing on MapReduce, *1st International Workshop on Big Data, Streams and Heterogeneous Source Mining (BigMine)*, 2012. (**Best Paper Award**)
32. L. Cao, Y.-C. Chang, N. Codella, M. Merler Q.-B. Nguyen, J. R. Smith, Multimedia Analytics: Modality Classification and Case-Based Retrieval tasks of ImageCLEF2012, *ImageCLEF Workshop 2012*. (**Winner of medical image classification track**)
31. N. Codella, A. Natsev, G. Hua, M. L. Hill, L. Cao, L. Gong, J. R. Smith, Video Event Detection Using Temporal Pyramids of Visual Semantics with Kernel Optimization and Model Subspace Boosting, *IEEE International Conference on Multimedia and Expo (ICME)*, 2012
30. Z. Li, H. Ning, L. Cao, T. Zhang, Y. Gong, T. S. Huang, Learning to Search Efficiently in High Dimensions, *Advances in Neural Information Processing Systems (NIPS)*, 2011.
29. Z. Yin, L. Cao, J. Han, C. Zhai, T. S. Huang, Geographical Topic Discovery and Comparison, *International Conference on World Wide Web (WWW)*, 2011. (Oral presentation)

28. S.-F. Tsai, L. Cao, F. Tang, T. S. Huang, Compositional Object Pattern: a New Model for Album Event Recognition, *ACM Conference on Multimedia (ACM MM)*, 2011.
27. Y. Lin, F. Lv, S. Zhu, M. Yang, T. Cour, K. Yu, L. Cao, T. S. Huang, Large-scale Image Classification: Fast Feature Extraction and SVM Training, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
26. Z. Yin, L. Cao, J. Han, C. Zhai, T. S. Huang, LPTA: A Probabilistic Model for Latent Periodic Topic Analysis, *IEEE International Conference on Data Mining (ICDM)*, 2011.
25. Z. Yin, L. Cao, J. Han, J. Luo, T. S. Huang, Diversified Trajectory Pattern Ranking in Geo-Tagged Social Media, *SIAM International Conference on Data Mining (SDM)*, 2011.
24. L. Cao, S.-F. Chang, N. Codella, C. Cotton, Dan. Ellis, L. Gong, M. Hill, G. Hua, J. Kender, M. Merler, Y. Mu, A. Natseve, J. Smith, IBM Research and Columbia University TRECVID-2011 Multimedia Event Detection (MED) System, *NIST TRECVID Workshop*, 2011.
23. L. Cao, Z. Liu, and T. S. Huang, Cross-dataset Action Detection, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
22. K. Yu, F. Lv, Y. Lin, S. Zhu, M. Yang, L. Cao, T. S. Huang, Large Scale Image Semantic Search, *European Conference on Computer Vision (ECCV) demo*, 2010.
21. L. Cao, A. Del Pozo, X. Jin, J. Luo, J. Han and T. S. Huang, RankCompete: Simultaneous Ranking and Clustering of Web Photos, *International Conference on World Wide Web (WWW)*, 2010.
20. J. Xin, A. Gallagher, L. Cao, J. Luo and J. Han, The Wisdom of Social Media: Using Flickr for Prediction and Forecast, *ACM Conference on Multimedia (ACM MM)*, 2010.
19. X. Jin, J. Han, L. Cao, J. Luo, B. Ding, C. Lin, Visual Cube and On-Line Analytical Processing of Images, *ACM Conference on Information and Knowledge Management (CIKM)*, 2010.
18. L. Cao, Y. Tian, Z. Liu, Z. Yao, Z. Zhang, T. S. Huang: Action detection using multiple spatial-temporal interest point features. *IEEE International Conference on Multimedia and Expo (ICME)*, 2010.
17. V. Le, H. Tang, L. Cao, T. Huang: Accurate and efficient reconstruction of 3D faces from stereo images. *IEEE International Conference on Image Processing (ICIP)* 2010.
16. L. Cao, J. Luo, A. Gallagher, X. Jin, J. Han, T. Huang: A worldwide tourism recommendation system based on geotaggedweb photos, *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* 2010.
15. X. Jin, S. Kim, J. Han, L. Cao, Z. Yin, GAD: General Activity Detection for Fast Clustering on Large Data, *SIAM International Conference on Data Mining (SDM)*, 2009.
14. Y. Hu, L. Cao, F. Lv, S. Yan, Y. Gong, and T. S. Huang, Event Detection in Complex Scenes with Spatial and Temporal Ambiguities, *IEEE International Conference on Computer Vision (ICCV)*, 2009. (Oral presentation)
13. L. Cao, J. Luo, F. Liang, and T. S. Huang, Heterogeneous Feature Machines for Visual Recognition, *IEEE International Conference on Computer Vision (ICCV)*, 2009.
12. L. Cao, J. Yu, J. Luo, and T. S. Huang, Enhancing Semantic and Geographic Annotation of Web Images via Logistic Canonical Correlation Regression, *ACM Conference on Multimedia (ACM MM)*, 2009.
11. L. Cao, J. Luo, T. S. Huang, Annotating Photo Collections by Label Propagation According to Multiple Similarity Cues, *ACM Conference on Multimedia (ACM MM)*, 2008.
10. Y. Fu, L. Cao, G. Guo, T. S. Huang: Multiple feature fusion by subspace learning. *ACM International Conference on Image and Video Retrieval (CIVR)*, 2008.
9. A. Gallagher, C. Neustaedter, J. Luo, L. Cao, T. Chen, Image Annotation Using Personal Calendars as Context, *ACM Conference on Multimedia (ACM MM)*, 2008.
8. L. Cao, M. Dikmen, Y. Fu, T. S. Huang, Gender Recognition from Body, *ACM Conference on Multimedia (ACM MM)*, 2008.

7. L. Cao, J. Luo, H. Kautz and T. S. Huang, Annotating Collections of Photos Using Hierarchical Event and Scene Models, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
6. L. Cao, and L. Fei-Fei, Spatially Coherent Latent Topic Model for Concurrent Segmentation and Classification of Objects and Scenes, *IEEE International Conference Computer Vision (ICCV)*, 2007.
5. S. Chen, L. Cao, J. Liu, X. Tang: Iterative MAP and ML Estimations for Image Segmentation, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
4. S. Chen, L. Cao, J. Liu, X. Tang: Automatic Segmentation of Lung Fields from Radiographic Images of SARS Patients Using a New Graph Cuts Algorithm. *International Conference on Pattern Recognition (ICPR)* 2006.
3. L. Cao, J. Liu, X. Tang: 3D object retrieval using 2D line drawing and graph based relevance reedback. *ACM Conference on Multimedia (ACM MM)*, 2006.
2. L. Cao, J. Liu and X. Tang, Degen Generalized Cylinders and Their Properties. *European Conference on Computer Vision (ECCV)*, 2006.
1. L. Cao, J. Liu and X. Tang. 3D Object Reconstruction from a Single 2D Line Drawing without Hidden Lines. *IEEE International Conference on Computer Vision (ICCV)*, 2005.

Patents

10. Y. Song, M. Gygli, L. Cao, Automatic video segment selection method and apparatus, US Patent 10,565,771 2020
9. L. Cao, J. Fan, C. Wang, B. Xiang, B. Zhou, Semi-supervised Learning for Word Embedding, US Patent 9,672,814, 2018.
8. L. Cao, W. Tan, L. Fong, Memory-aware Matrix Factorization, US Patent 9,947,073, 2018.
7. L. Cao, N. Duan, Q. Lin, C. Wang, J. Yan, and X. Zhang, Generating a Training Model Based on Feedback, US Patent 9,659,258, 2017.
6. L. Cao, Y. Chang, and Q. Nguyen, Association of Visual Labels and Event Context in Image Data, US Patent 9,734,166, 2017.
5. L. Cao, M. Merler, J. R. Smith, Systems and Methods for Inferring Gender by Fusion of Multimodal content, US Patent 9,471,851, 2016.
4. L. Cao, N. Codella, G. Hua, G. Leiguang, A. Natsev, and J. R. Smith, Techniques for Spatial Semantic Attribute Matching for Location Identification, US Patent 9,251,434, 2016.
3. Z. Liu and L. Cao, Adaptive Action Detection, US Patent 9,014,420, 2015.
2. L. Cao, N. Codella, G. Hua, J. R. Smith, Techniques for Ground-Level Photo Geolocation Using Digital Elevation, US Patent 9,165,217, 2015.
1. Z. Liu, Y. Tian, L. Cao, and Z. Zhang, Hierarchical filtered Motion Field for Action Recognition, US Patent 8,639,042, 2014.

RECENT PRESS COVERAGE

Forbes, Aug 2020

- Google Anthos Gains New Features Including Hybrid AI capabilities. ([link](#))

ZDNet, Aug 2020

- Google Cloud launches Anthos updates for identity, Kubernetes cluster, app modernization tools. ([link](#))

Google, July 2020

- Verizon Powers Intuitive Customer Experiences with Google Cloud. ([link](#))

ACM SIGMM, August 2017

- Dr. Liangliang Cao recives ACM SIGMM Rising Star Award 2017. ([link](#))

Techcrunch, April 2017

- ffVC and NYU are mixing researchers with investors in the name of machine intelligence. ([link](#))
This Week in Machine Learning&AI, April 2017
- Introducing the AI NexusLab Startups. ([link](#))
Techcrunch, Oct. 2016
- Detecting Sarcasm That Plays Off Images
Scientific American, Aug. 2016
- Computers Can Sense Sarcasm? Yeah, Right.
Yahoo Blog, July 2016
- Creating Animated GIFs Automatically from Video
NVidia DevBlog, May 2016
- Accelerate Recommender Systems with GPUs
VICE, May 2016
- Yahoo! Made an AI That Automatically Turns Videos Into Fire GIFs .
Taiwan CTS news, June 2016
- Yahoo Taiwan Deployed the Visual Search App for Online Shopping
Taiwan Nownews, Nov. 2015
- Yahoo! Supermarket App Powered Product Search with Image Queries
Inside Taiwan, Nov. 2015
- Yahoo! Supermarket Provides “Mobile Visual Search” App
Taiwan United Daily News, Nov. 2015
- Buy With Your Photo: A New Online Shopping App

CONFERENCE TUTORIALS

- | | |
|---|-----------------------------|
| Tutorial in ICME | July 2018, Shanghai, China |
| · Visual Search and Question Answering
Co-present with Lu Jiang, Yannis Kalant | |
| Tutorial in ACM Multimedia | Oct. 2013, Barcelona, Spain |
| · Massive-Scale Multimedia Semantic Modeling
Co-present with John R. Smith | |
| Tutorial in CVPR | June 2014, Columbus, OH |
| · Learning Visual Semantics: Models, Massive Computation and Applications
Co-present with Shih-Fu Chang, John R. Smith and Rogerio Feris | |

PROFESSIONAL SERVICES

Society

- IEEE, senior member
- ACM, senior member

Journal Editorship

- Associated Editor, IEEE Trans. Multimedia, 2020-now
- Associated Editor, IEEE Multimedia, 2019-now
- Associated Editor, Journal of Visual Communication and Image Representation, 2019-now

Conference Organizer

- Sponsor chair, ACM Conference on Multimedia (ACM MM), 2020
- Senior PC, AAI Conference on Artificial Intelligence, 2021

- Industrial chair, ACM International Conference on Multimedia Retrieval (ICMR), 2017
- Area chair, IEEE Winter Conference on Applications and Computer Vision (WACV) 2017
- Area chair, ACM Conference on Multimedia (ACM MM), 2017
- Workshop chair, ACM International Conference on Multimedia Retrieval (ICMR), 2016
- Area chair, IEEE Winter Conference on Applications and Computer Vision (WACV), 2014
- Area chair, ACM Conference on Multimedia (ACM MM), 2012
- Founding team, Greater New York Area Multimedia and Vision Meeting, 2012-2013
- Founder and chair, ACM Workshop on GeoMM, 2012-2015