

# LIANGLIANG CAO

<https://www.linkedin.com/in/liangliangcao/>

[www.llcao.net](http://www.llcao.net)

[liangliang.cao@gmail.com](mailto:liangliang.cao@gmail.com)

## WORK HISTORY

---

Principal Engineer / Director, Google DeepMind	09/2024 - now
Principal Scientist and Engineering Lead, Apple	10/2022 - 09/2024
Staff/Senior Staff Scientist and Senior Manager, Google	10/2018 - 10/2022
Co-founder and CTO, Switi Inc (acquired)	08/2016 - 10/2018
Senior Research Scientist, Yahoo Labs	04/2015 - 08/2016
Research Staff Member, IBM T. J. Watson Research Center	06/2011 - 04/2015

## EDUCATION AND ACADEMIA EXPERIENCE

---

Adjunct Research Professor, UMass Amherst	2018 - 2021
Adjunct Associate Professor, Columbia University in the City of New York	2013 - 2018
Ph.D., University of Illinois at Urbana-Champaign	2006 - 2011
M.Phil., Multimedia Lab, The Chinese University of Hong Kong	2003 - 2005
B. E., Dept. Electronic Engineering, University of Science and Technology of China	1998 - 2003

## SELECTED PROJECTS

---

- Gemini Live, Astra, and Gemini Multimodality: principal engineer/director, 2024-2025
- Apple Intelligence: modeling and engineering lead, 2023-2024
- Google Cloud Vision: tech lead to start the Cloud Multimodal Services, 2022
- Google Speech: team lead for Cloud speech modeling, 2018 - 2021
- Google-Verizon Contact Center AI: one of the largest contracts in the history of Cloud AI, 2019
- Switi Inc: Virtual Enterprise Representative Agent (Vera), 2016 - 2018
- ImageNet LSVRC Challenge: 1st place in 2010

## SELECTED AWARDS

---

- IEEE Fellow, Class 2025
- Google Cloud Feats of Engineering Award, 2019
- ACM SIGMM Rising Star Award, 2017
- IBM Outstanding Accomplishment for Multimedia Team, 2012
- IBM Watson Emerging Leader in Multimedia and Signal Processing, 2010
- First Place in ImageNet Large Scale Visual Recognition Challenge, 2010

## PUBLICATION

---

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

Total 1000+ citations, 15+ patents, H-Index 50.

Selected Recent Publications:

- X. Wei, M. Lin, F. Ye, F. Song, L. Cao, M. Thai, T. Yang, Learning with a Reference Model: From Generalization Theory to Scaling Law, ICML 2025
- D Xu, Y. Jiang, C. Huang, L. Song, T. Gernoth, L. Cao, Z. Wang, H. Tang, Cavia: Camera-controllable Multi-view Video Diffusion with View-Integrated Attention, ICML 2025

3. Y. Jiang, H. Tang, J.-H. Chang, L. Song, Z. Wang, L. Cao, Efficient-3DiM: Learning a Generalizable Single-image Novel-view Synthesizer in One Day, ICLR 2024
4. H. You, H. Zhang, Z. Gan, X. Du, B. Zhang, Z. Wang, L. Cao, S-F Chang, Y. Yang, Ferret: Refer and Ground Anything Anywhere at Any Granularity, ICLR 2024.
5. C. Chen, B. Zhang, L. Cao, J. Shen, T. Gunter, A. M.y Jose, A. Toshev, J. Shlens, R. Pang, Y. Yang, STAIR: Learning Sparse Text and Image Representation in Grounded Tokens, EMNLP 2023.
6. L. Cao, B. Zhang, C. Chen, Y. Yang, X. Du, W. Zhang, Z.n Lu, Y. Zheng Less is More: Removing Text-regions Improves CLIP Training Efficiency and Robustness, arXiv:2305.05095, 2023.
7. 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.
8. 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, IEEE Journal of Selected Topics in Signal Processing, 2021.
9. 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.