

Chenxi Liu

CONTACT INFORMATION	Email: cxliu@google.com Website: https://cs.jhu.edu/~cxliu/	
JOB	Google DeepMind , Mountain View, CA <i>Senior Research Scientist</i>	2023 -
	Waymo , Mountain View, CA <i>Senior Research Scientist</i> <i>Research Scientist</i>	2020 - 2023
EDUCATION	Johns Hopkins University Ph.D. in Computer Science Advisor: Alan Yuille Committee: Gregory Hager, Fei-Fei Li	2016 - 2020
	University of California, Los Angeles M.S. in Statistics Advisor: Alan Yuille	2015 - 2016
	Rice University Exchange Student in Electrical and Computer Engineering	Fall 2013
	Tsinghua University B.E. in Automation Minor in Economics	2011 - 2015
INTERNSHIP & VISITING	Facebook AI Research , Menlo Park, CA Research Intern Mentors: Saining Xie, Ross Girshick, Piotr Dollár, Kaiming He	Summer 2019
	Google , Sunnyvale, CA Software Engineering Intern Mentors: Fei-Fei Li, Wei Hua, Liang-Chieh Chen	Summer 2018
	Google , Sunnyvale, CA Software Engineering Intern Mentors: Jia Li, Wei Hua, Jonathan Huang, Jonathon Shlens, Barret Zoph, Kevin Murphy, Fei-Fei Li	Fall 2017
	Adobe Research , San Jose, CA Research Scientist Intern Mentors: Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu	Summer 2016
	Toyota Technological Institute at Chicago , Chicago, IL Research Assistant Advisors: Gregory Shakhnarovich, Michael Maire	Summer 2015
	University of Toronto , Toronto, ON Research Assistant Advisors: Raquel Urtasun, Sanja Fidler	Summer 2014

TEACHING

University of California, Santa Cruz

Fall 2021

Guest Lecturer

Course: Neural Computation

Instructor: Cihang Xie

Lecture: Trajectory Prediction: Towards Interactions and Joint Modeling

Johns Hopkins University

Fall 2018

Teaching Assistant

Course: Probabilistic Models of the Visual Cortex

Instructor: Alan Yuille

Lectures: Bayes Decision Theory, Vision and Language

PUBLICATIONS

Chen Wei, **Chenxi Liu**, Siyuan Qiao, Zhishuai Zhang, Alan Yuille, and Jiahui Yu. 2024. De-Diffusion Makes Text a Strong Cross-Modal Interface. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Seattle, Washington.

Tong He, Pei Sun, Zhaoqi Leng, **Chenxi Liu**, Dragomir Anguelov, and Mingxing Tan. 2023. LEF: Late-to-Early Temporal Fusion for LiDAR 3D Object Detection. In *International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, Detroit, Michigan.

Yingwei Li, Charles R. Qi, Yin Zhou, **Chenxi Liu**, and Dragomir Anguelov. 2023. MoDAR: Using Motion Forecasting for 3D Object Detection in Point Cloud Sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Vancouver, Canada, pages 9329-9339.

Zhaoqi Leng, Guowang Li, **Chenxi Liu**, Ekin Dogus Cubuk, Pei Sun, Tong He, Dragomir Anguelov, and Mingxing Tan. 2023. LidarAugment: Searching for Scalable 3D LiDAR Data Augmentations. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Computer Society, London, United Kingdom.

Pei Sun, Mingxing Tan, Weiyue Wang, **Chenxi Liu**, Fei Xia, Zhaoqi Leng, and Dragomir Anguelov. 2022. SWFormer: Sparse Window Transformer for 3D Object Detection in Point Clouds. In *Proceedings of European Conference on Computer Vision (ECCV)*. Springer, Tel Aviv, Israel.

Chenxi Liu, Zhaoqi Leng, Pei Sun, Shuyang Cheng, Charles R. Qi, Yin Zhou, Mingxing Tan, and Dragomir Anguelov. 2022. LidarNAS: Unifying and Searching Neural Architectures for 3D Point Clouds. In *Proceedings of European Conference on Computer Vision (ECCV)*. Springer, Tel Aviv, Israel.

Mao Ye, **Chenxi Liu**, Maoqing Yao, Weiyue Wang, Zhaoqi Leng, Charles R. Qi, and Dragomir Anguelov. 2022. Multi-Class 3D Object Detection with Single-Class Supervision. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Computer Society, Philadelphia, Pennsylvania, pages 5123-5130.

Zhaoqi Leng, Mingxing Tan, **Chenxi Liu**, Ekin Dogus Cubuk, Xiaojie Shi, Shuyang Cheng, and Dragomir Anguelov. 2022. PolyLoss: A Polynomial Expansion Perspective of Classification Loss Functions. In *International Conference on Learning Representations (ICLR)*.

Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, **Chenxi Liu**, Ashish Venugopal, David Weiss, Ben Sapp, Zhifeng Chen, and Jonathon Shlens. 2022. Scene Transformer: A unified architecture for predicting multiple agent trajectories. In *International Conference on Learning Representations (ICLR)*.

Scott Ettinger, Shuyang Cheng, Benjamin Caine, **Chenxi Liu**, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles Qi, Yin Zhou, Zoey Yang, Aurelien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. 2021. Large Scale Interactive Motion Forecasting for Autonomous Driving : The Waymo Open Motion Dataset. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Online, pages 9710-9719. **(Oral)**

Zefan Li, **Chenxi Liu**, Alan Yuille, Bingbing Ni, Wenjun Zhang, and Wen Gao. 2021. Progressive Stage-wise Learning for Unsupervised Feature Representation Enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Online, pages 9767-9776.

Alan Yuille and **Chenxi Liu**. 2021. Deep Nets: What have they ever done for Vision?. *International Journal of Computer Vision*. 129 (3), pages 781-802.

Chenxi Liu, Piotr Dollár, Kaiming He, Ross Girshick, Alan Yuille, and Saining Xie. 2020. Are Labels Necessary for Neural Architecture Search?. In *Proceedings of European Conference on Computer Vision (ECCV)*. Springer, Online, pages 798-813. **(Spotlight)**

Michelle Shu, **Chenxi Liu**, Weichao Qiu, and Alan Yuille. 2020. Identifying Model Weakness with Adversarial Examiner. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, New York, New York, pages 11998-12006.

Siyuan Qiao, Huiyu Wang, **Chenxi Liu**, Wei Shen, and Alan Yuille. 2019. Rethinking Normalization and Elimination Singularity in Neural Networks. *CoRR*, [abs/1911.09738](https://arxiv.org/abs/1911.09738).

Zhuotun Zhu, **Chenxi Liu**, Dong Yang, Alan Yuille, and Daguang Xu. 2019. V-NAS: Neural Architecture Search for Volumetric Medical Image Segmentation. In *Proceedings of the IEEE International Conference on 3D Vision (3DV)*. IEEE Computer Society, Quebec City, Canada, pages 240-248.

Siyuan Qiao, Huiyu Wang, **Chenxi Liu**, Wei Shen, and Alan Yuille. 2019. Weight Standardization. *CoRR*, [abs/1903.10520](https://arxiv.org/abs/1903.10520).

Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan Yuille, and Li Fei-Fei. 2019. Auto-DeepLab: Hierarchical Neural Architecture Search for Semantic Image Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Long Beach, California, pages 82-92. **(Oral)**

Xiaohui Zeng, **Chenxi Liu**, Yu-Siang Wang, Weichao Qiu, Lingxi Xie, Yu-Wing Tai, Chi Keung Tang, and Alan Yuille. 2019. Adversarial Attacks Beyond the Image Space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Long Beach, California, pages 4302-4311. **(Oral)**

Runtao Liu, **Chenxi Liu**, Yutong Bai, and Alan Yuille. 2019. CLEVR-Ref+: Diagnosing Visual Reasoning with Referring Expressions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Long Beach, California, pages 4185-4194.

Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. 2018. Progressive Neural Architecture Search. In *Proceedings of European Conference on Computer Vision (ECCV)*. Springer, Munich, Germany, pages 19-35. **(Oral)**

Siyuan Qiao, **Chenxi Liu**, Wei Shen, and Alan Yuille. 2018. Few-Shot Image Recognition by Predicting Parameters from Activations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Salt Lake City, Utah, pages 7229-7238. **(Spotlight)**

Yu-Siang Wang, **Chenxi Liu**, Xiaohui Zeng, and Alan Yuille. 2018. Scene Graph Parsing as Dependency Parsing. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. Association for Computational Linguistics, New Orleans, Louisiana, pages 397-407. **(Oral)**

Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. 2017. Recurrent Multimodal Interaction for Referring Image Segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Venice, Italy, pages 1280-1289.

Yan Wang, Lingxi Xie, **Chenxi Liu**, Siyuan Qiao, Ya Zhang, Wenjun Zhang, Qi Tian, and Alan Yuille. 2017. SORT: Second-Order Response Transform for Visual Recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Venice, Italy, pages 1368-1377.

Siyuan Qiao, Wei Shen, Weichao Qiu, **Chenxi Liu**, and Alan Yuille. 2017. ScaleNet: Guiding Object Proposal Generation in Supermarkets and Beyond. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Venice, Italy, pages 1809-1818.

Chenxi Liu, Junhua Mao, Fei Sha, and Alan Yuille. 2017. Attention Correctness in Neural Image Captioning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, San Francisco, California, pages 4176-4182.

Chenxi Liu*, Alexander Schwing*, Kaustav Kundu, Raquel Urtasun, and Sanja Fidler. 2015. Rent3D: Floor-Plan Priors for Monocular Layout Estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Boston, Massachusetts, pages 3413-3421. **(Oral)**

Jingwen Bai, **Chenxi Liu**, and Ashutosh Sabharwal. 2014. Increasing Cellular Capacity Using ISM Band Side-channels: A First Study. In *Proceedings of the 4th workshop on All things cellular: operations, applications, & challenges*. ACM, Chicago, Illinois, pages 9-14.

AWARDS

2nd Place, Waymo Perception Ping Pong Championship	2023
Outstanding Reviewer, CVPR	2021
MINDS Doctoral Dissertation Award	2020
Outstanding Reviewer, ECCV	2020

2019 Google PhD Fellowship	2019
Finalist, 2019 Adobe Research Fellowship	2018
Finalist, 2018 NVIDIA Graduate Fellowship	2018
2nd Place, Google Sunnyvale Moffett Place Pool Tournament	2017
Honorable Mention, 2017 Snap Research Fellowship	2017
Finalist, 2018 Adobe Research Fellowship	2017
Beijing Outstanding Graduate	2015
Tsinghua University Excellent Graduate	2015
National Southwest Associated University Scholarship, Tsinghua University	2014
Meritorious Winner, Mathematical Contest in Modeling, COMAP	2014
Geru Zheng Scholarship, 1st class, Tsinghua University	2013
National Scholarship, Tsinghua University	2013
3rd Prize, Beijing Universities Physics Competition	2012
EMC Scholarship, Tsinghua University	2012
Freshman Scholarship, 2nd class, Tsinghua University	2011

INVITED TALKS

On the Automation and Diagnosis of Visual Intelligence	
– MINDS Symposium, Johns Hopkins University	October 2020
Neural Architecture Search: Acceleration and Generalization	
– Facebook	January 2020
– ByteDance	February 2020
– Amazon	February 2020
– Waymo	February 2020
– Facebook	March 2020
Progressive Neural Architecture Search	
– Google Research Conference	October 2018
– Johns Hopkins University	October 2018
– Leiphone Webinar	September 2018
– European Conference on Computer Vision	September 2018
– Stanford University (Host: Fei-Fei Li, Amir Zamir)	July 2018
– Tsinghua University (Host: Jiwen Lu)	April 2018
– Valse Webinar (Host: Wei Shen)	December 2017
Rent3D: Floor-Plan Priors for Monocular Layout Estimation	
– IEEE Conference on Computer Vision and Pattern Recognition	June 2015

PATENTS

Wei Hua, Barret Zoph, Jonathon Shlens, **Chenxi Liu**, Jonathan Huang, Jia Li, Fei-Fei Li, Kevin Murphy. Neural architecture search using a performance prediction neural network. *US Patent 16861491*.

Xin Lu, Zhe Lin, Xiaohui Shen, Jimei Yang, Jianming Zhang, Jen-Chan Jeff Chien, **Chenxi Liu**. Deep salient content neural networks for efficient digital object segmentation. *US Patent 10460214*.

Zhe Lin, Xin Lu, Xiaohui Shen, Jimei Yang, **Chenxi Liu**. Automatically segmenting images based on natural language phrases. *US Patent 10089742*.

Chenxi Liu, Tianlin Shi, Xinyi Yang. Method for identifying and correcting piano rhythms by using intelligent robot. *CN Patent 101930680*.

SERVICE

Co-organizer of LVVU@CVPR 2020.

Invited reviewer for NIPS 2016, CVPR 2019, JMLR, NAACL 2019, PAMI, ICCV 2019, Neurocomputing, ACL 2019, EMNLP 2019, TMM, AAAI 2020, CVPR 2020, ACL 2020, ECCV 2020, NAS@ICLR 2020, EMNLP 2020, AACL-IJCNLP 2020, AutoML@ICML 2020, BMVC 2020, IJCV, CVPR 2021, ICCV 2021, ICLR 2022, CVPR 2022, AutoML 2022, ECCV 2022, NeurIPS 2022, ICLR 2023.

Reviewer for CVPR 2017, ICIP 2017, ICCV 2017, NIPS 2017, AAAI 2018, PRCV 2018, Pattern Recognition, AAAI 2019, ICML 2019.

MISCELLANEOUS

Programming Languages	Python, Matlab, C++, C, \LaTeX
Deep Learning Tools	PyTorch, Tensorflow, Caffe, Theano
TOEFL	117 (Reading 29, Listening 30, Speaking 29, Writing 29)
GRE	335 (Verbal 160, Quantitative 170, Writing 5.0)