

HANG ZHANG

+1(732) 543 · 4857

zhanghang0704@gmail.com

ABOUT ME

I am Hang Zhang, a Senior Staff Applied Research Scientist at Cruise, leading the efforts in Detection, Segmentation and Perception model consolidation. Before joining Cruise, I was a Research Scientist at Meta, and lead the effort in building a generic and scalable architecture optimization platform for AI products, serving various production models for IG, Portal and VR headsets on person understanding, AR/VR rendering and Ads ranking. Before joining Meta, I was a Senior Applied Scientist in Amazon AI, where I worked on computer vision, deep learning and MXNet framework. We built ResNeSt model which achieved state-of-the-arts results on several major computer vision tasks.

Beyond my work, I am also enthusiastic in contributing to open source projects, including [D2Go](#), [Detectron2](#), [AutoGluon](#), [PyTorch Encoding](#) and [GluonCV](#). More about me:

[\[Homepage\]](#) [\[GitHub\]](#) [\[LinkedIn\]](#) [\[Google Scholar\]](#)

EXPERIENCE

Cruise AI

Senior Staff Applied Research Scientist

Sep 2022 - Now

San Francisco, CA

- Lead the efforts in camera major model consolidation on object detection and segmentation in 2023. Leading the efforts in building a multi-task, multi-view, multi-modality, multi-frame and multi-platform model at Perception.
- Lead the development of the Perception model training pipeline, which supports various projects on detection & segmentation, longtailed understanding, lane detection, offboard foundation model, and online distillation.

Meta Reality Lab (Facebook)

Research Scientist

Oct 2020 - Aug 2022

Menlo Park, CA

- Lead the development of FBNAS project, a unified pipeline for cross-platform hardware-aware model optimization. FBNAS has been applied to several production models in person understanding on IG, AR/VR applications and Ads models.
- Developed and open sourced D2Go toolkit, bringing Detectron2 to mobile [\[post\]](#)
- Research on efficient architectures, e.g. [FBNetV5](#), [ScaleViT](#). Co-organized workshop on “Computer Vision for MetaVerse” in ECCV2022.

Amazon AWS AI

Senior Applied Scientist

Jan 2018 - Oct 2020

East Palo Alto, CA

- Lead the development of GluonCV toolkit [\[link\]](#) and AutoGluon toolkit (for AutoML) [\[link\]](#).
- Lead research on large scale vision solution, e.g. [ResNeSt](#), [Bag-of-tricks](#), [CFNet](#), [dynamic SGD](#).
- Organized 3 tutorials on [ICCV19](#), [CVPR20](#) and [ECCV20](#).

Amazon Lab 126 (Internship)

Applied Scientist Intern

May 2017 - Aug 2017

Cupertino, CA

- Developed SoTA semantic segmentation algorithm of [EncNet](#) (Oral paper (~2.1%) in CVPR 2018) [\[Talk link on YouTube\]](#)

NVIDIA (Internship)

Deep learning Research Intern

May 2016 - Aug 2016

Holmdel, NJ

- Developed an end-to-end deep learning solution for autonomous driving.
- Implemented Torch to Caffe model converter [[GitHub](#)].

EDUCATION

Rutgers University

2013 - 2017

Ph.D. in Electrical and Computer Engineering

Thesis Advisor: Prof. Kristin Dana

Research Interest: Computer Vision

Current GPA: 3.9/4.0

Southeast University (Nanjing, China)

2009 - 2013

B.S. in School of Automation

Advisor: Prof. Junyang Li

Outstanding Undergraduate Thesis 2013 - School of Automation, Southeast University

PUBLICATIONS

1. Runjian Chen, **Hang Zhang**, Avinash Ravichandran, Wenqi Shao, Alex Wong, and Ping Luo. Clap: Unsupervised 3d representation learning for fusion 3d perception via curvature sampling and prototype learning. *In Submission to CVPR*, 2025
2. Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, **Hang Zhang**, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023
3. Bichen Wu, Chaojian Li, **Hang Zhang**, Xiaoliang Dai, Peizhao Zhang, Matthew Yu, Jialiang Wang, Yingyan Lin, and Peter Vajda. Fbnetv5: Neural architecture search for multiple tasks in one run. *arXiv preprint arXiv:2111.10007*, 2021
4. Chaojian Li, Kyungmin Kim, Bichen Wu, Peizhao Zhang, **Zhang, Hang**, Xiaoliang Dai, Peter Vajda, and Yingyan Lin. An investigation on hardware-aware vision transformer scaling. *ACM Transactions on Embedded Computing Systems*, 23(3):1–19, 2024
5. **Zhang, Hang**, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. ResNeSt: Split-Attention Networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2736–2746, 2022
6. Yi Zhu, Zhongyue Zhang, Chongruo Wu, Zhi Zhang, Tong He, **Hang Zhang**, R Manmatha, Mu Li, and Alexander Smola. Improving semantic segmentation via efficient self-training. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021
7. **Hang Zhang**, Han Zhang, Chenguang Wang, and Junyuan Xie. Co-occurrent features in semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019
8. Tong He, Zhi Zhang, **Hang Zhang**, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks to train convolutional neural networks for image classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019
9. Haibin Lin, **Hang Zhang**, Yifei Ma, Zhi Zhang, Sheng Zha, and Mu Li. Elastic distributed training: Learning in the limbo of resources. *arXiv preprint arXiv:1904.12043*, 2019
10. Jian Guo, He He, Tong He, Leonard Lausen, Mu Li, Haibin Lin, Xingjian Shi, Chenguang Wang, Junyuan Xie, Sheng Zha, Aston Zhang, **Hang Zhang**, Zhi Zhang, Zhongyue Zhang, and Shuai Zheng. Gluoncv and gluonnlp: Deep learning in computer vision and natural language processing. *arXiv preprint arXiv:1907.04433*, 2019

11. Parneet Kaur, **Hang Zhang**, and Kristin Dana. Photo-realistic facial texture transfer. In *Winter Conference on Applications of Computer Vision (WACV)*, 2019
12. **Hang Zhang**, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Ambrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018 (**oral**)
13. Jia Xue, **Hang Zhang**, and Kristin Dana. Deep texture manifold for ground terrain recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018
14. **Hang Zhang** and Kristin Dana. Multi-style generative network for real-time transfer. *European Conference of Computer Vision Workshops (ECCVW)*, 2018
15. **Hang Zhang**. *Reflectance and texture encoding for material recognition and synthesis*. PhD thesis, Rutgers University-School of Graduate Studies, 2017
16. **Hang Zhang**, Jia Xue, and Kristin Dana. Deep ten: Texture encoding network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017
17. Jia Xue, **Hang Zhang**, Kristin Dana, and Ko Nishino. Differential angular imaging for material recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017
18. **Hang Zhang**, Kristin Dana, and Ko Nishino. Friction from reflectance: Deep reflectance codes for predicting physical surface properties from one-shot in-field reflectance. In *European Conference on Computer Vision (ECCV)*, pages 808–824. Springer, 2016
19. **Hang Zhang**, Kristin Dana, and Ko Nishino. Reflectance hashing for material recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3071–3080, 2015

ACADEMIA SERVICES

Workshop and Tutorial Organizer

<i>European Conference on Computer Vision (ECCV)</i>	Tel Aviv, 2022
Computer Vision for MetaVerse	

<i>European Conference on Computer Vision (ECCV)</i>	Glasgow, 2020
From HPO to NAS: Automatic Deep Learning.	

<i>IEEE Conference on Computer Vision and Pattern Recognition (CVPR)</i>	Seattle, 2020
From HPO to NAS: Hands-on Tutorial on Automatic Deep Learning.	

<i>IEEE International Conference on Computer Vision (ICCV)</i>	Seoul, 2019
Everything You Need to Know to Reproduce SOTA Deep Learning Models: Hands-on Tutorial for Training SOTA Computer Vision Models.	

<i>Amazon Machine Learning Conference (AMLC)</i>	Seattle, 2018
CNNs for Semantic Segmentation.	

Reviewer for Journals:

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
IEEE Transactions on Biomedical Circuits and Systems (TbioCAS)
Computer Vision and Image Understanding (CVIU)

Program Committee and Reviewer for Conferences:

<i>IEEE Conference on Computer Vision and Pattern Recognition (CVPR)</i>	2018 - 2021
--	-------------

<i>IEEE International Conference on Computer Vision (ICCV)</i>	2019 - 2021
<i>European Conference on Computer Vision (ECCV)</i>	2018 - 2020
<i>Conference on Neural Information Processing Systems (NeurIPS)</i>	2020 - 2021
<i>IEEE Winter Conference on Applications of Computer Vision (WACV)</i>	2018 - 2019
SIGGRAPH	2018