

ZHANGJIE CAO

Stanford University, Gates Computer Science Building, Room 156

☎ 650-441-4422 | ✉ caozj@cs.stanford.edu | 🏠 caozhangjie.github.io | 📱 caozhangjie

Education

Stanford University

PH.D. OF COMPUTER SCIENCE

U.S.A

September 2018 - Now

Tsinghua University

BACHELOR OF SOFTWARE ENGINEERING

• GPA: 91/100

China

September 2014 - July 2018

Research Interests

Robotics, Reinforcement Learning, Imitation Learning, Transfer Learning, Machine Learning, Deep Learning, Computer Vision

Honors & Awards

- | | | |
|------|--|---------------------|
| 2015 | National Scholarship (rank 1/64) , Tsinghua University | China |
| 2015 | Best Project Award (rank 1/56) , Course of Architecture of Computer and Network | China |
| 2016 | Qualcomm Scholarship , Tsinghua University | China |
| 2018 | Rank 3 in Visual Domain Adaptation Challenge (VisDA-2018) , ECCV2018 Workshop Challenge (http://ai.bu.edu/visda-2018) | Munich, Germany |
| 2019 | Outstanding reviewers , CVPR 2019 | Long Beach, CA, USA |

Publications

- Zhangjie Cao**, Yilun Hao, Mengxi Li, Dorsa Sadigh. **Learning Feasibility to Imitate Demonstrators with Different Dynamics**. *Conference on Robot Learning (CoRL)*, 2021.
- Bo Fu*, **Zhangjie Cao***, Jianmin Wang, Mingsheng Long. **Transferable Query Selection for Active Domain Adaptation**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. (Oral)
- Yang Shu*, **Zhangjie Cao***, Chenyu Wang, Jianmin Wang, Mingsheng Long. **Open Domain Generalization with Domain-Augmented Meta-Learning**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Chao Huang*, **Zhangjie Cao***, Yunbo Wang*, Jianmin Wang, Mingsheng Long. **MetaSets: Meta-Learning on Point Sets for Generalizable Representations**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Zhangjie Cao**, Dorsa Sadigh. **Learning from Imperfect Demonstrations from Agents with Varying Dynamics**. *Robotics and Automation Letters (RA-L) with International Conference on Robotics and Automation (ICRA)*, 2021.
- Zhangjie Cao***, Minae Kwon*, Dorsa Sadigh. **Transfer Reinforcement Learning across Homotopy Classes**. *Robotics and Automation Letters (RA-L) with International Conference on Robotics and Automation (ICRA)*, 2021.
- Zhangjie Cao***, Erdem Biyik*, Woodrow Z. Wang, Allan Raventos, Adrien Gaidon, Guy Rosman, Dorsa Sadigh. **Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving**. *Robotics: Science and Systems (RSS)*, 2020.
- Amir Zamir, Alexander Sax, Jitendra Malik, Nikhil Cheerla, Rohan Suri, **Zhangjie Cao**, Leonidas Guibas. **Robust Learning Through Cross-Task Consistency**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

9. Kaidi Cao, Jingwei Ji*, **Zhangjie Cao***, Chien-Yi Chang, Juan Carlos Niebles. **Few-Shot Video Classification via Ordered Temporal Alignment**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
10. Bingbin Liu, Ehsan Adeli, **Zhangjie Cao**, Kuan-Hui Lee, Abhijeet Shenoj, Adrien Gaidon, Juan Carlos Niebles. **Spatiotemporal Relationship Reasoning for Pedestrian Intent Prediction**. *International Conference on Robotics and Automation (ICRA)*, 2020.
11. Yuxuan Song, Lantao Yu, **Zhangjie Cao**, Zhiming Zhou, Jian Shen, Shuo Shao, Weinan Zhang, Yong Yu. **Improving Unsupervised Domain Adaptation with Variational Information Bottleneck**. *European Conference on Artificial Intelligence (ECAI)*, 2020.
12. Boxiao Pan, **Zhangjie Cao***, Ehsan Adeli, Juan Carlos Niebles. **Adversarial Cross-Domain Action Recognition with Co-Attention**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
13. Aditya Grover, Christopher Chute, Rui Shu, **Zhangjie Cao**, Stefano Ermon. **AlignFlow: Learning from Multiple Domains via Normalizing Flows**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
14. **Zhangjie Cao***, Kaichao You*, Mingsheng Long, Jianmin Wang, Qiang Yang. **Learning to Transfer Examples for Partial Domain Adaptation**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
15. Hong Liu*, **Zhangjie Cao***, Mingsheng Long, Jianmin Wang, Qiang Yang. **Separate to Adapt: Open Set Domain Adaptation via Progressive Separation**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
16. Kaichao You, Mingsheng Long, **Zhangjie Cao**, Jianmin Wang, Michael I. Jordan. **Universal Domain Adaptation**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
17. Yang Shu, **Zhangjie Cao**, Mingsheng Long, Jianmin Wang. **Transferable Curriculum for Weakly-Supervised Domain Adaptation**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
18. Mingsheng Long, **Zhangjie Cao**, Jianmin Wang, Michael I. Jordan. **Conditional Adversarial Domain Adaptation**. *Neural Information Processing Systems (NIPS)* 2018.
19. **Zhangjie Cao**, Lijia Ma, Mingsheng Long, Jianmin Wang. **Partial Adversarial Domain Adaptation**. *European Conference on Computer Vision (ECCV)*, 2018.
20. Mingsheng Long, **Zhangjie Cao**, Jianmin Wang, Han Zhu, Michael I. Jordan. **Learning Transferable Visual Features with Very Deep Adaptation Networks**. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
21. **Zhangjie Cao**, Mingsheng Long, Ziping Sun, Jianmin Wang. **Deep Priority Hashing**. *ACM Multimedia Conference (ACM MM)*, 2018.
22. **Zhangjie Cao**, Mingsheng Long, Jianmin Wang, Michael I. Jordan. **Partial Transfer Learning with Selective Adversarial Networks**. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
23. **Zhangjie Cao**, Mingsheng Long, Chao Huang, Jianmin Wang. **Transfer Adversarial Hashing for Hamming Space Retrieval**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
24. Zhongyi Pei*, **Zhangjie Cao***, Mingsheng Long, Jianmin Wang. **Multi-Adversarial Domain Adaptation**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
25. Mingsheng Long, **Zhangjie Cao**, Jianmin Wang, Philip S. Yu. **Learning Multiple Tasks with Multilinear Relationship Networks**. *Neural Information Processing Systems (NIPS)*, 2017.
26. **Zhangjie Cao**, Mingsheng Long, Jianmin Wang, Philip S. Yu. **HashNet: Deep Learning to Hash by Continuation**. *International Conference on Computer Vision (ICCV)*, 2017.
27. **Zhangjie Cao**, Qixing Huang, Ramani Karthik. **3D Object Classification via Spherical Projections**. *International Conference on 3D Vision (3DV)*, 2017.

28. **Zhangjie Cao**, Mingsheng Long, Qiang Yang. **Transitive Hashing Network for Heterogeneous Multimedia Retrieval**. *AAAI Conference on Artificial Intelligence (AAAI)*, 2017. (Oral)

Professional Experience

Reviewer for Conferences, Journals

Computer Vision: CVPR 2019-2021, ICCV 2019-2021, ECCV 2020, BMVC 2020, WACV 2021, IJCV 2019

Machine Learning: ICML 2019,2021, NeurIPS 2019-2021, ICLR 2020-2021

Artificial Intelligence: AAAI 2020-2022, IJCAI 2021-2022

Robotics: CoRL 2020-2021, ICRA/RA-L 2021-2022

Work Experience

Waymo LLC (self-driving company)

U.S.A

Mentor: Yin Zhou

REALISTIC SCENE SIMULATION

June. 2019 - Sept. 2019

- Work as a research intern on a project for realistic simulation of full scene to address simulation difficulty on rarely existing corner cases and novel objects, which has huge impact on Waymo.
- Propose a novel generative model to refine the fake input scene and generate realistic scene with the same semantic as the input.
- Writing the main part of the paper for the work.

Research Experience

Department of Computer Science, Stanford University

U.S.A

Mentor: Dorsa Sadigh

PROJECT: LEARNING FROM IMPERFECT DEMONSTRATIONS

Sept. 2020 - Now

- Raise the problem of standard imitation learning that the assumption of the same dynamics and optimal demonstrations can be violated.
- Propose a imitation learning setting where the demonstrations could be collected from different dynamics or be suboptimal.
- Design several methods to address learning from demonstrations of different dynamics or suboptimal demonstrations.
- The works are published in ICRA 2021, CoRL 2021.

Department of Computer Science, Stanford University

U.S.A

Mentor: Dorsa Sadigh

PROJECT: HUMAN MODELING IN NEAR-ACCIDENT SCENARIOS

Sept. 2019 - Jan. 2019

- Propose a method to model human driving behavior in near accident scenarios.
- Outperform imitation learning and hard-coded policies in aspects of comfort and safety.
- The work is published in RSS 2020.

Department of Computer Science, Stanford University

U.S.A

Mentor: Amir R. Zamir

PROJECT: CROSS-TASK CONSISTENCY

May 2019 - June 2019

- Proposed a flexible and fully computational framework for learning with consistency across an arbitrary dictionary of tasks.
- Proved that consistency constraints can improve the generalization power of the model to out-of-distribution data.
- The paper has been published in CVPR 2020.

Department of Computer Science, Stanford University

U.S.A

Mentor: Stefano Ermon

PAPER: ALIGNFLOW: CYCLE CONSISTENT LEARNING FROM MULTIPLE DOMAINS VIA NORMALIZING FLOWS

Feb. 2019 - May 2019

- Proposed to use normalizing flows in generative model, which is flexible in loss (Adversarial training or Maximum likelihood).
- AlignFlow guarantees exact cycle consistency in mapping datapoints from one domain to another.
- The paper has been published in AAAI 2020.

Department of Computer Science, Stanford University

U.S.A

Mentor: Juan Carlos Niebles

PROJECT: LEARNING FOR VIDEOS

Oct. 2018 - Feb. 2019

- Address intent prediction by modeling the relation between pedestrians and environment objects with spatiotemporal graph.
- The method achieved good results on real-world driving datasets and the paper has been published in ICRA 2020.
- Proposed cross-domain action recognition to recognize actions of an unlabeled domain with a labeled related domain.
- Designed a co-attention module and a temporal-aligned adaptation module for global action distribution matching.
- The method achieved significant performance and the paper has been published at AAAI 2020.

Department of Computer Science, The University of Texas at Austin

U.S.A

Mentor: Qixing Huang

PAPER: 3D OBJECT CLASSIFICATION VIA SPHERICAL PROJECTIONS

Feb. 2017 - May. 2017

- Proposed a spherical representation leveraging depth variation and contour information for 3D objects.
- Developed deep neural networks composing of two parts for depth and contour representation respectively to classify 3D objects.
- Implemented Spherical Projection in caffe framework and carefully designed experiments to compare our method with state of the art methods under standard evaluation criteria on large scale 3D Recognition Dataset.
- Wrote the first version of the paper under the supervision of my mentor.

National Lab for Big Data Systems, School of Software, Tsinghua University

China

Mentor: Mingsheng Long

PROJECTS ON DOMAIN ADAPTATION

Sept. 2016 - June. 2018

- Proposed a Multi-Adversarial Domain Adaptation to train multiple adversaries weighted by probability over classes.
- Proposed a new partial domain adaptation setting with source label space includes target and a selective adversarial network with multiple adversarial networks and both instance-level and class-level weights to address it.
- Further proposed a single adversarial network architecture to address partial domain adaptation

National Lab for Big Data Systems, School of Software, Tsinghua University

China

Mentor: Mingsheng Long

PROJECTS ON DEEP LEARNING TO HASH

Jan. 2016 - Nov. 2017

- Proposed a new cross-modal retrieval scenario without explicit relationship and a Transitive Hashing Network to solve it.
- Proposed HashNet with binary output and enable optimization of sign activation function by continuation method.
- Proposed Transfer Adversarial Hashing the first model focusing on cross-domain retrieval within Hamming Radius 2.
- Proposed Deep Focal Hashing (DFH) with priority loss to address the class imbalance and easy-hard imbalance problems.

Department of Computer Science and Lewis-Sigler Institute of Integrative

U.S.A

Genomics, Princeton University

Mentor: Olga Troyanskaya

PROJECT: DEEP LEARNING FOR NONCODING VARIANTS RECOGNITION

July. 2016 - August. 2016

- Designed new architecture with multiple classifiers build on each convolutional layers to exploit low level features and solved the gradient vanishing problem of the original network.
- Considering the relation of different chromatin features (labels), I improved the classifying layer with low rank technical.
- The new architecture outperformed existing methods under standard evaluation criteria such as AUC of PR Curve and ROC Curve.