

Yingwei Li

CONTACT INFORMATION

Department of Computer Science
Johns Hopkins University

yingwei.li@jhu.edu
<http://yingwei.li>

RESEARCH INTERESTS

My research interests mainly lay in computer vision and deep learning, especially on autonomous driving [13,14,15,19], robust representation learning [3,4,7,8,10,13,14,16,17,18,19], multi-modality fusion [2,13,18,19], neural architecture search [5,6,9,11], medical machine intelligence [1,2,3]. Representative papers are **highlighted in bold**. I am open to new topics.

EDUCATION

Johns Hopkins University (JHU) 09/2018 - 07/2022 (expected)
Ph.D. in Computer Science
Advisor: Alan Yuille

National Taiwan University (NTU) Spring 2017
Exchange Student in Computer Science and Information Engineering

Fudan University (Fudan) 2014 - 2018
B.S. in Computer Science, *Honor Class*

EXPERIENCE

Google, Research Intern 06/2021 - present
Working on accurate and robust multi-modality 3D object detection.

- Mentors: Dr. Mingxing Tan, Dr. Denny Zhou, Mr. Jiquan Ngiam and Dr. Adams Wei Yu
- Improved **20.7% relative performance** of internal single-modality baseline, which achieves the state-of-the-art results on Waymo Open Dataset.
- Improved **6.3% relative performance** and achieved **strong robustness** against corrupted and out-of-distribution data based on the state-of-the-art single-modality model by the proposed cross-attention based multi-modality fusion module. Submitted one paper to CVPR [19] and plan to file one US patent.

Waymo, Software Engineering Intern 05/2020 - 11/2020
Worked on accurate and robust multi-modality long-range object distance estimation.

- Mentors: Prof. Hang Zhao, Dr. Ruichi Yu, Dr. Maya Kabkab and Dr. Tiffany Yu-Han Chen
- **Built long-range distance estimation datasets and the state-of-the-art distance estimation pipeline.**
- Reduced **13.3% relative estimation error** and achieved **strong out-of-distribution robustness** based on the state-of-the-art distance estimation baseline. Filed one US patent and submitted one paper to ICLR [13].

ByteDance AI Lab, Research Intern 05/2019 - 11/2019
Worked on Neural Architecture Search (NAS) and lightweight deep learning model design.

- Mentors: Dr. Linjie Yang, Dr. Xiaojie Jin and Dr. Xiaochen Lian
- **Built an NAS pipeline from scratch using the ByteDance infrastructure** and the searched lightweight models achieved the state-of-the-art performance on ImageNet.
- Proposed an lightweight attention module that improved **1.1% top-1 ImageNet accuracy** compared with the prior state-of-the-art model. Filed one CN patent and published one paper on CVPR [6].

Johns Hopkins University, Research Assistant

09/2018 - 07/2022 (expected)

Working on assessing and improving the deep learning model robustness.

- Advisor: Prof. Alan Yuille
- Proposed a debiased neural network training strategy that improves **the model accuracy and different kinds of model robustness**. Published one paper on ICLR [10].
- Analyzed the threat of adversarial attacks on the state-of-the-art volumetric medical image segmentation model and show how to defense against the attack. Published a book chapter [3].
- Proposed a model augmentation strategy to improve the transferability of adversarial examples. Published one paper on AAAI [4].

ByteDance AI Lab, University Collaboration Program

09/2018 - 05/2019

Worked on generating transferable universal adversarial perturbations that fool defense models.

- Mentor: Dr. Xiaohui Shen
- Found a simple but powerful universal perturbation can fool a series of state-of-the-art defenses. Published one paper on ECCV [7].

TuSimple, Research Intern

06/2016 - 09/2016

Worked on efficient multiple object tracking under autonomous driving scenario.

- Mentor: Dr. Naiyan Wang
- Accelerated the tracking algorithm by simultaneously extracting the appearance feature and detecting the object with a shared backbone network. The proposed method reduced nearly 50% executing time with 2% detection accuracy improvement.

PREPRINTS

Representative papers are **highlighted in bold**; *: equally contribution.

- [19] **Yingwei Li***, Adams Wei Yu*, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc Le, Alan Yuille, Mingxing Tan. **Lidar-Camera Deep Fusion for Multi-Modal 3D Object Detection**. Under Review.
- [18] Junfei Xiao, Longlong Jing, Lin Zhang, Ju He, Qi She, Zongwei Zhou, Alan Yuille, **Yingwei Li**. Learning from Temporal Gradient for Semi-supervised Action Recognition. Under Review.
- [17] Vipul Gupta, Zhuowan Li, Adam Kortylewski, Chenyu Zhang, **Yingwei Li**, Alan Yuille. Swap-Mix: Diagnosing and Regularizing the Over-reliance on Visual Context in Visual Question Answering. Under Review.
- [16] Shunchang Liu, Jiakai Wang, Aishan Liu, **Yingwei Li**, Yijie Gao, Xianglong Liu, Dacheng Tao. Harnessing Perceptual Adversarial Patches for Crowd Counting. Under Review.
- [15] Longlong Jing, Ruichi Yu, Jiyang Gao, Henrik Kretzschmar, Kang Li, Charles R. Qi, Hang Zhao, Alper Ayvaci, Xu Chen, Dillon Cower, **Yingwei Li**, Yurong You, Han Deng, Congcong Li, Dragomir Anguelov. Depth Matters Most: Improving Per-Object Depth Estimation for Monocular 3D Detection and Tracking. Under Review.
- [14] Ziqi Zhang, Xinge Zhu, **Yingwei Li**, Xiangqun Chen, Yao Guo. Adversarial Attacks on Monocular Depth Estimation. In *CoRR, abs/2003.10315*.

- [13] **Yingwei Li**, Tiffany Chen, Maya Kabkab, Ruichi Yu, Longlong Jing, Yurong You, Hang Zhao. **R4D: Utilizing Reference Objects for Long-Range Distance Estimation**. In *International Conference on Learning Representations (ICLR)*, 2022.
- [12] Jieru Mei, Yucheng Han, Yutong Bai, Yixiao Zhang, **Yingwei Li**, Xianhang Li, Alan Yuille, Cihang Xie. Fast AdvProp. In *International Conference on Learning Representations (ICLR)*, 2022.
- [11] Huaijin Pi, Huiyu Wang, **Yingwei Li**, Zizhang Li, Alan Yuille. Searching for TrioNet: Combining Convolution with Local and Global Self-Attention. In Proceedings of *the British Machine Vision Conference (BMVC)*, BMVA Press, 2021.
- [10] **Yingwei Li**, Qihang Yu, Mingxing Tan, Jieru Mei, Peng Tang, Wei Shen, Alan Yuille, Cihang Xie. **Shape-Texture Debaised Neural Network Training**. In *International Conference on Learning Representations (ICLR)*, 2021.
- [9] Qihang Yu, **Yingwei Li**, Jieru Mei, Yuyin Zhou, Alan L. Yuille. CAKES: Channel-wise Automatic Kernel Shrinking for Efficient 3D Network. In Proceedings of *the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2021.
- [8] Song Bai, **Yingwei Li**, Yuyin Zhou, Qizhu Li, Philip H.S. Torr. Adversarial Metric Attack for Person Re-identification. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, IEEE, 2020.
- [7] **Yingwei Li**, Song Bai, Cihang Xie, Zhenyu Liao, Xiaohui Shen, Alan Yuille. Regional Homogeneity: Towards Learning Transferable Universal Adversarial Perturbations Against Defenses. In Proceedings of *the European Conference on Computer Vision (ECCV)*, Springer, 2020.
- [6] **Yingwei Li**, Xiaojie Jin, Jieru Mei, Xiaochen Lian, Linjie Yang, Cihang Xie, Qihang Yu, Yuyin Zhou, Song Bai, Alan Yuille. **Neural Architecture Search for Lightweight Non-Local Networks**. In Proceedings of *the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2020.
- [5] Jieru Mei, **Yingwei Li**, Xiaochen Lian, Xiaojie Jin, Linjie Yang, Alan Yuille, Jianchao Yang. AtomNAS: Fine-Grained End-to-End Neural Architecture Search. In *International Conference on Learning Representations (ICLR)*, 2020.
- [4] **Yingwei Li**, Song Bai, Yuyin Zhou, Cihang Xie, Zhishuai Zhang, Alan Yuille. Learning Transferable Adversarial Examples via Ghost Networks. In Proceedings of *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, AAAI Press, 2020.
- [3] **Yingwei Li***, Zhuotun Zhu*, Yuyin Zhou, Yingda Xia, Wei Shen, Elliot K. Fishman, and Alan L. Yuille. Volumetric Medical Image Segmentation: A 3D Deep Coarse-to-fine Framework and Its Adversarial Examples. In *Deep Learning and Convolutional Neural Networks for Medical Image Computing*, Advances in Computer Vision and Pattern Recognition, Springer, ISBN 978-3-030-13968-1, 2019.
- [2] Yuyin Zhou, **Yingwei Li**, Zhishuai Zhang, Yan Wang, Angtian Wang, Elliot K. Fishman, Alan Yuille, Seyoun Park. Hyper-Pairing Network for Multi-Phase Pancreatic Ductal Adenocarcinoma Segmentation. In Proceedings of *the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, 2019.
- [1] Yuyin Zhou, David Dreizin, **Yingwei Li**, Zhishuai Zhang, Yan Wang, Alan Yuille. Multi-Scale Attentional Network for Multi-Focal Segmentation of Active Bleed after Pelvic Fractures. In Proceedings of *10th International Workshop on Machine Learning in Medical Imaging (MLMI)* Held in Conjunction with MICCAI, Springer, 2019.

TALKS	Lidar-Camera Deep Fusion for Multi-Modal 3D Object Detection	
	– Google Brain Waymo meeting	Dec, 2021
	– Google Cloud Vision/Video Tech Talk	Jan, 2022
	R4D: Utilizing Reference Objects for Long-Range Distance Estimation	
	– Google Brain Waymo meeting	July, 2021
	Shape-Texture Debiased Neural Network Training	
	– Qingyuan Seminars	Feb, 2021
	– Visual Informatics Group @ University of Texas at Austin	Sep, 2021
	Learning Transferable Adversarial Examples via Ghost Networks	
	– AdvML Workshop @ CVPR 2019	June, 2019
– The Thirty-Fourth AAAI Conference on Artificial Intelligence	Feb, 2020	
Neural Architecture Search for Lightweight Non-Local Networks		
– Kwai Silicon Valley Lab	May, 2020	
SELECTED AWARDS	ICLR Travel Award	2020
	First Prize Scholarship from Fudan University Education Development Foundation	2017
	SCSK Scholarship	2016
	Silver Medal, ACM-ICPC Shanghai Regional Contest	2014
	Bronze Medal, China National Olympiad in Informatics (NOI)	2013
	First Prize, China National Olympiad in Informatics in Provinces (NOIP)	2012 & 2013
SERVICE	Co-organizer of	
	– The Art of Robustness: Devil and Angel in Adversarial Machine Learning	CVPR 2022
	– Practical Deep Learning in the Wild	AAAI 2022
	– Adversarial Robustness in the Real World	ICCV 2021
	– Adversarial Learning for Multimedia	ACMMM 2021
	– Adversarial Robustness in the Real World	ECCV 2020
	Reviewer for IEEE TIP, IEEE TDSC, Neurocomputing, Pattern Recognition, AmlCV@CVPR2020, SRML@ICML2021, SecMI@ICLR2021, RseMI@AAAI2021 AAAI 2021, IJCAI 2021, CVPR 2021, ICCV 2021, NeurIPS 2021, AAAI 2022, ICLR 2022, CVPR 2022.	
TEACHING	Johns Hopkins University (JHU)	Spring 2021
	Role: Teaching Assistant	
	Course: EN.601.783 <i>Vision as Bayesian Inference</i>	
	Instructor: Alan Yuille	
ADVISING	Junfei Xiao	Master student from Johns Hopkins University
	Weiyu Guo	Master student from University of Chinese Academy of Sciences
	Shunchang Liu	Undergraduate student from Beihang University
SKILLS	Python, TensorFlow and PyTorch (for research projects); C/C++ (for ACM-ICPC contests).	