CHENGZHI MAO

Ph.D. Candidate, Columbia University
+1 (917)459-8209 ⋄ mcz@cs.columbia.edu
http://www.cs.columbia.edu/˜mcz/

INTEREST

Computer Vision, Robust Machine Learning, Foundation Model

EDUCATION

Columbia University *Aug 2018 - 2023(Expected)*
Ph.D. in Computer Science
Co-advised by Prof. Carl Vondrick and Prof. Junfeng Yang

Tsinghua University *Aug 2013 - Jul 2018*
B.S., Electronic Engineering
Advised by Prof. Yuan Shen

Massachusetts Institute of Technology *Jun 2017 - Dec 2017*
Visiting Student at CSAIL
Advised by Prof. Dina Katabi

PUBLICATIONS

1. Doubly Right Object Recognition: A Why Prompt for Visual Rationales. *CVPR2023.* **Chengzhi Mao**, Revant Teotia, Amrutha Sundar, Sachit Menon, Junfeng Yang, Xin Wang, Carl Vondrick.
2. Understanding Zero-shot Adversarial Robustness for Large-Scale Models. *ICLR 2023.* **Chengzhi Mao**, Scott Geng, Junfeng Yang, Xin Wang, Carl Vondrick.
3. Shape Analysis by Shadow Synthesis. *CVPR 2023.* Ruoshi Liu, Sachit Menon, **Chengzhi Mao**, Dennis Park, Simon Stent, Carl Vondrick.
4. Causal Transportability for Visual Recognition. *CVPR 2022.* **Chengzhi Mao**, Kevin Xia, James Wang, Hao Wang, Junfeng Yang, Elias Bareinboim, Carl Vondrick.
5. Discrete Representations Strengthen Vision Transformer Robustness. *ICLR 2022.* **Chengzhi Mao**, Lu Jiang, Mostafa Dehghani, Carl Vondrick, Rahul Sukthankar, Irfan Essa.
6. Real-Time Neural Voice Camouflage. *ICLR 2022. (Oral, top 2%)* Mia Chiquier, **Chengzhi Mao**, Carl Vondrick.
7. Adversarial Attacks are Reversible with Natural Supervision. *ICCV 2021.* **Chengzhi Mao**, Mia Chiquer, Hao Wang, Junfeng Yang, Carl Vondrick.
8. Generative Interventions for Causal Learning. *CVPR 2021.* **Chengzhi Mao**, Amogh Gupta, Augustine Cha, Hao Wang, Junfeng Yang, Carl Vondrick.
9. Multitask Learning Strengthens Adversarial Robustness. *ECCV 2020. (Oral, top 2%).* **Chengzhi Mao**, Amogh Gupta, Vikram Nitin, Baishakhi Ray, Shuran Song, Junfeng Yang, Carl Vondrick.
10. Metric Learning for Adversarial Robustness. *NeurIPS 2019.* **Chengzhi Mao**, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, Baishakhi Ray.
11. Bidirectional Inference Networks: A Class of Bayesian Networks for Health Profiling. *AAAI 2019* Hao Wang, **Chengzhi Mao**, Hao He, Mingmin Zhao, Tommi Jaakkola, Dina Katabi

12. A Tale of Two Models: Constructing Evasive Attadks on Edge Models. Proceedings of Machine Learning and System. Wei Hao, Asaf Cidon, Junfeng Yang, Aahil Awatramani, **Chengzhi Mao**, Jiayang Hu, Pin-Chun Chen, Eyal Cidon.
13. A Probabilistic Learning Approach to UWB Ranging Error Mitigation. *IEEE GLOBECOM 2018.* **Chengzhi Mao**, Kangbo Lin, Tiancheng Yu, Yuan Shen.

TECHNICAL REPORTS

1. Self-Supervised Convolutional Visual Prompts. *Submitted to ICML.* Yun-Yun Tsai*, **Chengzhi Mao***, Yow-Kuan Lin, Junfeng Yang.
2. Robustifying Language Models with Test-Time Adaptation. *ICLRW 2023.* Noah McDermott, Junfeng Yang, **Chengzhi Mao**.
3. Robust Perception through Equivariance. *Submitted to ICML.* **Chengzhi Mao**, Lingyu Zhang, Abhishek Joshi, Junfeng Yang, Hao Wang, Carl Vondrick.
4. Adversarially Robust Video Perception by Seeing Motion. *Submitted to ICCV.* **Chengzhi Mao***, Lingyu Zhang*, Junfeng Yang, Carl Vondrick.
5. Landscape Learning for Optimization-Based Inference. *Submitted to ICCV.* Ruoshi Liu, **Chengzhi Mao**, Purva Tendulkar, Hao Wang, Carl Vondrick.
6. Test-time Defense against Adversarial Attacks: Detection and Reconstruction of Adversarial Examples via Masked Autoencoder. *CVPRW 2023.* Yun-Yun Tsai, Ju-Chin Chao, Albert Wen, Zhaoyuan Yang, **Chengzhi Mao**, Tapan Shah, Junfeng Yang
7. Natural-Parameter Networks as a Class of General-Purpose Probabilistic Neural Networks. Hao Wang, **Chengzhi Mao**, Xingjian Shi, Dit-Yan Yeung.
8. Live Trojan Attacks on Deep Neural Networks. *CVPR 2020 Adversarial Workshop.* Robby Costales, **Chengzhi Mao**, Raphael Norwitz, Bryan Kim, Junfeng Yang.
9. Fooling Semantic Segmentation in One Step via Manipulating Nuisance Factors. *ECCV 2020 Adversarial Learning Workshop.* Guangyu Shen, **Chengzhi Mao**, Junfeng Yang, Baishakhi Ray
10. Using Multiple Self-Supervised Tasks Improves Model Robustness. *ICLR 2022 Workshop.* Matthew Lawhon, **Chengzhi Mao**, Junfeng Yang.

WORK EXPERIENCE

Microsoft Research. Robust Foundation Models. Mentor: Xin Wang. 2022 Jun-Sep.

Google Research. Robust Vision Transformer model. Mentor: Lu Jiang, Rahul Sukthankar, Mostafa Dehghani, and Irfan Essa. 2021 Jun-Dec.

Waymo. Multitask learning for autonomous driving's perception model. Mentor: Paul Donnelly and Chen Zhao. 2020 Jun-Sep.

TEACHING EXPERIENCE

Columbia University *Jan 2022 - Jun 2022*
Head Teaching Assistant for 450 students
Developed course material and assignments for **Computer Vision II**.
Instructor: Carl Vondrick

Columbia University *Jan 2019 - Jun 2019*
Head Teaching Assistant
Developed course material and give lectures for **Security and Robustness of Machine Learning**.
Instructor: Junfeng Yang

STUDENTS MENTORSHIP

Scott Geng	Zero-Shot Robustness	May 2022-Now
Revant Teotia	Doubly Right Recognition	Sep 2021-Now
Lingyu Zhang	Robust Inference from Motion	Jan 2021-Now
Amrutha Varshini Sundar	Doubly Right Recognition	Sep 2022-Now
Noah Thomas MaDermott	Robust NLP via Test Time Adaptation	Jan 2022-Now
Abhishek Joshi	Robustness via Equivariance	Jan 2022-Now
Matthew Lawhon	Multitask Robustness	Sep 2021-May 2022
James Wang	Causal Computer Vision	Jan 2021-May 2022
Cynthia Mao	Test Time Inference via Generative Model	Sep 2019-Jan 2020
Yu Li	Adversarial Training on Low-Rank Purified Images	Sep 2019-Jan 2020
Augustine Cha	Steering Generative Models	Sep 2019-Jun 2020
Amogh Gupta	Multitask Learning	Sep 2019-Jun 2020
Guangyu Shen	Generating Realistic Image Attacks	May 2019-Oct 2019
Robby Costales	Trojan Attack for Neural Network	May 2019-Oct 2019
Ziyuan Zhong	Adversarial Robust Visual Classifier	Jan 2019-May 2019

SKILLS

Framework: Pytorch, Tensorflow, Jax
Programming Language: Python, C++, C, MATLAB, Lua, SQL, Verilog, LaTeX, HTML

SERVICES

Journal/Conference Reviewer: TPAMI 2021-2022, ICLR 2020-2023, NeurIPS 2020-2022, AAAI 2021-2023, CVPR 2020-2023, ICML 2021-2022, ICCV 2021, ECCV 2022, WACV 2022, BMVC 2022.

Community Service: Give two talks on "Demystifying the PhD" at Columbia University to help students from diversified background to learn about PhD, Serve as a Graduate Application Advisor for underrepresented groups

TALKS

McGill University	Reliable Machine Learning by Integrating Context	March 9, 2023
Wabbi	Reliable Machine Learning by Integrating Context	Feb 22, 2023
MIT	Reliable Machine Learning by Integrating Context	Feb 14, 2023
RPI	Reliable Machine Learning by Integrating Context	Feb 6, 2023
Tsinghua University	Reliable Machine Learning by Integrating Context	Feb 1, 2023
Rutgers University	Reliable Machine Learning by Integrating Context	Jan 10, 2023
Hong Kong University	Reliable Machine Learning by Integrating Context	Jan 6, 2023