YUZHE MA

University of Wisconsin–Madison
ma234@wisc.edu · 6082133287

Education

University of Wisconsin–Madison *09/2016-present*
Ph.D. in Computer Sciences, Minor in Statistics
Advisor: Professor Xiaojin (Jerry) Zhu

University of Wisconsin-Madison *05/2018*
M.S. in Computer Sciences

Huazhong University of Science and Technology *06/2016*
B.S. in Computer Science and Technology

Research Interests

My research interest lies in machine learning, with a focus on analyzing the adversarial robustness of models in sequential decision making problems, including multi-armed bandit and reinforcement learning. I am also interested in solving real-world problems related to trust-worthy artificial intelligence, unsupervised learning (e.g., dimensionality reduction), deep learning and differential privacy.

Work Experience

Research Intern, IBM Research *06-09/2020*

- Built a two-stage machine learning model based on Cycle Generative Adversarial Network (Cycle-GAN) and object detection techniques to identify buildings on the satellite imagery data of American cities.
- Developed an iterative training procedure based on the object detection algorithm YOLO to augment the labels in the training data. The model increased the total amount of labeled buildings from 20K to 80K.
- Applied Cycle-GAN to transform imagery data into rasterized maps.

Applied Scientist Intern, Amazon *06-09/2019*

- Developed a student identification model using the Gradient Boosted Tree (GBT) algorithm.
- Carried out an end-to-end machine learning pipeline, including data acquisition, model training, hyper-parameter tuning, post-processing of predictions, and model testing.
- Evaluated the model performance on Amazon Prime data and achieved 85% accuracy.

Research Intern, Symantec Research Labs (NortonLifeLock) *05-08/2018*

- Proposed a federated machine teaching framework that coordinates the training process of local nodes to jointly teach some desired model, while respecting the local data privacy during the communication between nodes.
- The work is published in the International Joint Conference on Neural Networks (**IJCNN**).

Research Scholar, Cornell University *07-08/2015*

- Proposed a nonlinear dimensionality reduction algorithm, which is able to preserve both the global and the local structure of high-dimensional data after reduction.
- The work is published in the International Frontiers of Algorithmics Workshop (**FAW**), and the extended version appeared in the Theoretical Computer Science (**TCS**).
- Advised by Prof. Kun He & Prof. John E. Hopcroft.

Publication

Superscript ⋆ for alphabetic author order.

Yun-Shiuan Chuang, Xuezhou Zhang, **Yuzhe Ma**, Mark K. Ho, Joseph L. Austerweil, Xiaojin Zhu. Using Machine Teaching to Investigate Human Assumptions when Teaching Reinforcement Learners. In The 43rd Annual Meeting of the Cognitive Science Society (**CogSci**), 2021.

Yuzhe Ma, Jon Sharp, Ruizhe Wang, Earlence Fernandes, Xiaojin Zhu. Sequential Attacks on Kalman Filter-based Forward Collision Warning Systems. In The NDSS Automotive and Autonomous Vehicle Security Workshop (**NDSS-AutoSec**), 2021. (Demo paper)

Yuzhe Ma, Jon Sharp, Ruizhe Wang, Earlence Fernandes, Xiaojin Zhu. Sequential Attacks on Kalman Filter-based Forward Collision Warning Systems. In The 35th AAAI Conference on Artificial Intelligence (**AAAI**), 2021.

Xuezhou Zhang, Shubham Bharti, **Yuzhe Ma**, Adish Singla, Xiaojin Zhu. The Sample Complexity of Teaching by Reinforcement on Q-learning. In The 35th AAAI Conference on Artificial Intelligence (**AAAI**), 2021.

Xuezhou Zhang, **Yuzhe Ma**, Adish Singla. Task-agnostic Exploration in Reinforcement Learning. In The 34th Conference on Neural Information Processing Systems (**NeurIPS**), 2020.

Xuezhou Zhang, **Yuzhe Ma**, Adish Singla, Xiaojin Zhu. Adaptive Reward-Poisoning Attacks against Reinforcement Learning. In The 37th International Conference on Machine Learning (**ICML**), 2020.

Yuzhe Ma, Xuezhou Zhang, Wen Sun, Xiaojin Zhu. Policy Poisoning in Batch Reinforcement Learning and Control. In The 33rd Conference on Neural Information Processing Systems (**NeurIPS**), 2019.

Yuzhe Ma, Xiaojin Zhu, Justin Hsu. Data Poisoning against Differentially-Private Learners: Attacks and Defenses. In The 28th International Joint Conference on Artificial Intelligence **(IJCAI**), 2019.

Yufei Han, **Yuzhe Ma**, Chris Gates, Kevin Roundy, Yun Shen. Collaborative and Privacy-Preserving Machine Teaching via Consensus Optimization. In International Joint Conference on Neural Networks (**IJCNN**), 2019.

Kwang-Sung Jun⋆, Lihong Li⋆, **Yuzhe Ma**⋆, Xiaojin Zhu⋆. Adversarial Attacks on Stochastic Bandits. In The 32nd Conference on Neural Information Processing Systems (**NeurIPS**), 2018.

Yuzhe Ma, Kwang-Sung Jun, Lihong Li, Xiaojin Zhu. Data Poisoning Attacks in Contextual Bandits. In The 9th Conference on Decision and Game Theory for Security (**GameSec**), 2018.

Ayon Sen, Scott Alfeld, Xuezhou Zhang, Ara Vartanian, **Yuzhe Ma**, Xiaojin Zhu. Training Set Camouflage. In The 9th Conference on Decision and Game Theory for Security (**GameSec**), 2018.

Yuzhe Ma, Robert Nowak, Philippe Rigollet, Xuezhou Zhang, Xiaojin Zhu. Teacher Improves Learning by Selecting a Training Subset. In The 21st International Conference on Artificial Intelligence and Statistics (**AISTATS**), 2018.

Yuzhe Ma, Kun He, John Hopcroft, Pan Shi. Neighbourhood-Preserving Dimension Reduction via Localised Multidimensional Scaling. In Theoretical Computer Science (**TCS**), 2017.

Yuzhe Ma, Kun He, John Hopcroft, Pan Shi. Nonlinear Dimension Reduction by Local Multidimensional Scaling. In The 10th International Frontiers of Algorithmics Workshop (**FAW**), 2016.

Yuzhe Ma, Kun He, Leihua Qin, Yan Wang. A Primary Research on Overlapping Community Detection. In The 32nd National Conference on Theoretical Computer Science (**NCTCS**), 2014.

Patents

Yufei Han, **Yuzhe Ma**, Chris Gates, Kevin Roundy, Yun Shen. Systems and Methods for Preventing Decentralized Malware Attacks, U.S. Patent, 11,025,666, 2021.

Academic Service

Program Committee: ACML21, AAAI21, ACML20, AAAI20, ACML19, AAAI19
Conference Reviewer: ICLR22, NeurIPS21, ICML21, AISTATS21, ICLR21, NeurIPS20, ICML20, AISTATS20, NeurIPS19, ICML19, AISTATS19
Journal Reviewer: TON, TPAMI, IEEE Access, Machine Learning
Student Volunteer: AISTATS18

Honors and Awards

Student Travel Award, NeurIPS	*2019*
Top 50% Reviewer, NeurIPS	*2019*
Student Travel Award, GameSec	*2018*
Honorarium Award, GameSec Special Track	*2018*
Student Travel Award, AISTATS	*2018*
UW CS Summer Research Award	*2017*
Outstanding Bachelor Thesis Award	*2016*
CCF Outstanding Undergraduate Award	*2015*
Outstanding Student Leader Award	*2014*
China National Scholarship Award	*2013*

Skills & Expertise

Programming Skills:	Python, Pytorch, SQL, Matlab, C, C++, R, AMPL, Verilog
Machine Learning:	Multi-Armed Bandit, Reinforcement Learning, Recommender System, Deep Learning