

Tong Wu

9015 Eager Road #194, Richmond Heights, MO 63144
tongwu@wustl.edu | 765-720-4989 | <https://tongwu2020.github.io/tongwu/>

RESEARCH INTEREST

(Trustworthy) Machine Learning, Security, and Computer Vision

EDUCATION

Washington University in St. Louis – *GPA: 4.0/4.0* St. Louis, MO
Bachelor/Master of Science; Computer Science Major; Mathematics Major; Aug 2018 - May 2021
Relevant Coursework: Adversarial Artificial Intelligence (A, graduate), Computer Vision (A, graduate), Bayesian Machine Learning (A, graduate), Applications of Deep Neural Networks (A+, graduate), Multivariate Statistical Analysis (A, graduate)

DePauw University, College of Liberal Arts – *GPA: 3.94/4.0 (Major GPA: 4.0/4.0)* Greencastle, IN
Bachelor of Arts; Pre-Engineering Major; Mathematics Minor; Aug 2016 - May 2018
Relevant Coursework: Statistical Computing (A), Data Structures (A), Object-Oriented Software Development (A)

PUBLICATIONS

- **Tong Wu**, Liang Tong, and Yevgeniy Vorobeychik. “Defending Against Physically Realizable Attacks on Image Classification” In Proceedings of the 8th International Conference on Learning Representations (ICLR), May 2020. (Spotlight, acceptance rate 6.01%) URL: <https://arxiv.org/abs/1909.09552>
- Shaojie Wang, **Tong Wu**, and Yevgeniy Vorobeychik, “Towards Robust Sensor Fusion in Visual Perception” (Preprint) URL: <https://arxiv.org/abs/2006.13192>

RESEARCH EXPERIENCE

Defending against Physically Realizable Attacks on Image Classification (ICLR 2020, Spotlight) St. Louis, MO
Research Intern supervised by Prof. Yevgeniy Vorobeychik Dec 2018 - Dec 2019

- Studied the problem of defending deep neural network approaches for image classification from physically realizable attacks
- Demonstrated that the state-of-the-art robust models exhibit limited effectiveness against three highest profile physical attacks
- Proposed a new abstract model, ROA, in which an adversary places a small crafted rectangle that fools the image classifier
- Adversarial Training using our ROA achieved much better robustness against physically realizable attacks than all SOTA

Towards Robust Sensor Fusion in Visual Perception St. Louis, MO
Research Intern at TRustworthy Autonomous Systems Engineering Lab (TRASE) Dec 2019 - May 2020
Supervised by Prof. Yevgeniy Vorobeychik, Prof. Ayan Chakrabarti

- Illustrated the robustness of sensor fusion models against image-only and LiDAR-only attacks by exhaustive experiments
- Developed gradient-based camera-and-LiDAR combined attacks, which proved the fusion methods are also vulnerable

Optical Trojans: Assisting Adversarial Perturbations with Coded Defocus St. Louis, MO
Research Intern at TRustworthy Autonomous Systems Engineering Lab Aug 2020 - Present
Supervised by Prof. Yevgeniy Vorobeychik, Prof. Ayan Chakrabarti, and Prof. Xuan Zhang

- Designed optical lens which can assist the adversarial perturbations via coded defocus while maintaining the natural accuracy
- Demonstrated that such lens could be easily deployed in real world by evaluating the performance under various lens' positions, quantization constraints and noise inside lens

Robustness of Speaker Recognition and Identification Remote due to COVID-19
Research Intern at Cleverhans Lab supervised by Prof. Nicolas Papernot May 2020 - Aug 2020

- Illustrated the performance degradation of adversarial attacks reconstructed from spectrogram to audio via Griffin-Lim and true-phase inverse short time Fourier transform algorithms
- Investigated the property that Griffin-Lim algorithm could implicitly modify the phase information of the given spectrogram

TEACHING EXPERIENCE

Washington University in St. Louis

St. Louis, MO

Teaching Assistant of Introduction to Machine Learning

Jan 2019 - May 2020

- Collaborated with Professor to lead all teaching assistants on determining and evaluating the rubrics for assignments
- Held regular office hour every week, helped students on course materials especially theoretical analysis of machine learning
- Graded students' lab assignments and exams; primarily checking the mathematical proof and coding efficiency and functionality

SECELECTED PROJECT

Mitigating the Adversarial Behaviors in Crowdsourcing

St. Louis, MO

*Course Project of Human-in-the-Loop Computation supervised by **Prof. Chien-Ju Ho***

Sep 2020 - Present

- Simulated backdoor attacks in the crowdsourcing and model training process by designing triggers under various threat models
- Designed a robust pipeline by adaptive task assignment and semi-supervised learning to mitigate the adversarial effects

SELECTED HONORS

Washington University Graduate Affiliation Scholarship

2019, 2020, 2021

Washington University Undergraduate Research Conference Travel Award

2020

DePauw University Merit Scholarship

2016, 2017, 2018

DePauw Dean's List

2016, 2017, 2018

Putnam Mathematical Competition top 10%

2018

Michigan Competition MATH Challenge 3/74

2018

PROFESSIONAL ACTIVITES

Program Committee of AAAI 2021

2020

Shadow Program Committee of IEEE S&P 2021

2020

Volunteer of ICLR 2020 & ICML 2020

2020

Member of Tau Beta Pi Association

2019, 2020

DePauw Science Research Fellow

2017, 2018

Judge of West Central Indiana Regional Science and Engineering Fair

2018

SKILLS

Programming Languages: Python (Proficient), R (Proficient), MATLAB (Proficient), C++, Java, and Mathematica

DL Framework & Other Techniques: Pytorch, TensorFlow, Keras, Scikit-Learn, Numpy, Pandas, OpenCV, and Linux

Languages: Mandarin (Native), English (Proficient)

Others: Participant of Chinese Go Game association, Volunteer of Beijing Museum of Natural History