Jihye Choi

jihye@cs.wisc.edu
https://jihyechoi77.github.io

Interests

Make the practice of machine learning (ML) to be more explainable and more reliable.

Education

University of Wisconsin-Madison, Madison, WI — Since Aug 2019
Ph.D., Computer Sciences (Advisor: Prof. Somesh Jha)

- Relevant Coursework: Numerical Optimization, Computational Learning Theory, Verified Deep Learning, Big Data Systems, Mathematical Foundations of Machine Learning, Human-Computer Interaction, Advanced Algorithms

Carnegie Mellon University, Pittsburgh, PA — Aug 2016 - Dec 2017
M.S., Electrical and Computer Engineering

Yonsei University, Seoul, Korea — Mar. 2013 - Feb. 2016
B.S., School of Integrated Technology

- Early graduation with Highest Honors and Distinction, College of Engineering

Publications

CADE: Concept-based Adaptive Out-of-Distribution Detection and Explanation
Jihye Choi, Jayaram Raghuram, Somesh Jha
In Preparation

Rethink Diversity in Deep Learning Testing
Zi Wang, Jihye Choi, Somesh Jha
arXiv 2023

Why Train More? Effective and Efficient Membership Inference via Memorization
Jihye Choi, Varun Chandrasekaran, Shruti Tople, Somesh Jha
arXiv 2023

Identifying and Mitigating the Security Risks of Generative AI
Clark Barrett, Brad Boyd, Ellie Burzstein, Nicholas Carlini, Brad Chen, Jihye Choi, Amrita Roy Chowdhury, Mihai Christodorescu, Anupam Datta, Soheil Feizi, Kathleen Fisher, Tatsunori Hashimoto, Dan Hendrycks, Somesh Jha, Daniel Kang, Florian Kerschbaum, Eric Mitchell, John Mitchell, Zulfikar Ramzan, Khawaja Shams, Dawn Song, Ankur Taly, Diyi Yang
Workshop on Securing the Future of GenAI: Mitigating Security Risks, arXiv 2023

Concept-based Explanations for Out-Of-Distribution Detectors
Jihye Choi, Jayaram Raghuram, Ryan Feng, Jiefeng Chen, Somesh Jha, Atul Prakash
ICML 2023

Stratified Adversarial Robustness with Rejection
Jiefeng Chen*, Jayaram Raghuram*, Jihye Choi, Xi Wu, Yingyu Liang, Somesh Jha
(* equal contribution)
ICML 2023

Revisiting Adversarial Robustness of Classifiers With a Reject Option

Jiefeng Chen*, Jayaram Raghuram*, Jihye Choi, Xi Wu, Yingyu Liang, Somesh Jha
AAAI 2022 Workshop (**Oral Presentation and Best Paper Award**)

Stochastic Doubly Robust Gradient
Kanghoon Lee*, Jihye Choi*, Moonsoo Cha, Jung-Kwon Lee and Tae Yoon Kim
arXiv 2018

Data-driven Approach to Aesthetic Enhancement
Jihye Choi, Sungjoon Koh, Jongwoo Kwack, Yonghun Kwon, Hyunjung Shim
SPIE Electronic Imaging 2016

Work Experience

University of Wisconsin-Madison — Madison, WI
Research Assistant with Prof. Somesh Jha — Since Jun 2020

Visa Research — Palo Alto, CA
PhD Intern in Identity and Authentication team — May 2023 - Aug 2023

- Secure and robust federated learning with reduced computational overhead
- Mentored by Dr. Rahul Rachuri, Dr. Ke Wang, Dr. Yizhen Wang

Visa Research — Palo Alto, CA
PhD Intern in System Security Team — May 2022 - Aug 2022

- Robustifying ML models for source codes against semantic-preserving adversarial transformations.
- Mentored by Dr. Ke Wang, Dr. Yizhen Wang

CyLab, Carnegie Mellon University — Pittsburgh, PA
Research Assistant with Prof. Lujo Bauer — May 2017 - Mar 2019

- Investigate security and privacy threats in modern machine learning.
- Develop deep neural network architectures for face recognition and verification systems to make them more robust to evasion attacks involving adversarial examples.

T-Brain, SK Telecom — Seoul, Korea
Research Intern — Apr 2018 - Sep 2018

- Develop a stochastic gradient optimization method with double robustness to efficiently mitigate biases in missing data and reduce the variance in stochastic gradient descent.

Vision & Learning Lab., Yonsei University — Seoul, Korea
Undergraduate Research Assistant with Prof. Hyunjung Shim — Mar 2015 - Aug 2016

- Propose an automated image editing framework that improves the overall aesthetic quality of photos involving contextual modifications.

Teaching

UW-Madison CS 435 Intro to Cryptography — Fall 2021, Fall 2020, Spring 2020
(taught by Prof. Somesh Jha)

UW-Madison CS 412 Numerical Analysis — Fall 2019
(taught by Prof. Amos Ron)

CMU 18-290 Signals and Systems — Spring 2017

(taught by Prof. Aswin C Sankaranarayanan and Prof. Richard Stern)

Skills

Programming: Python (incl. Keras, Tensorflow, Jax, PyTorch), C, C++, MATLAB, R, Spark, SQL
Languages: English (Advanced), Korean (Native)

References

Available upon request