# Haolin Yuan

781-290-9017 | hyuan4@jh.edu

## Education

**Johns Hopkins University** — Baltimore, MD
*Doctor of Philosophy in Computer Science* — *Jan. 2022 – now*

**Johns Hopkins University** — Baltimore, MD
*Master of Science in Security Informatics* — *Aug. 2019 – Dec. 2020*

**Brandeis University** — Waltham, MA
*Bachelor of Science in Computer Science*
*Bachelor of Art in Mathematics* — *Aug. 2014 – May 2018*

## Publication

**SneakyPrompt: Jailbreaking Text-to-image Generative Models**
Yuchen Yang, Bo Hui, **Haolin Yuan**, Neil Gong, and Yinzhi Cao
*in the Proceedings of the IEEE Symposium on Security and Privacy (Oakland), 2024*

**EdgeMixup: Embarrassingly Simple Data Alteration to Improve Lyme Disease Lesion Segmentation and Diagnosis Fairness**
**Haolin Yuan**, John Aucott, Armin Hadzic, William Pual, Marcia Villegas de Flores, Phillip Mathew, Phillip Burlina, Yinzhi Cao.
*in the Proceedings of the 26th Medical Image Computing and Computer Assisted Intervention (MICCAI), 2023*

**PrivateFL: Accurate, Differentially Private Federated Learning via Personalized Data Transformation**
**Haolin Yuan***, Bo Hui*, Yuchen Yang*, Neil Gong, Yinzhi Cao.
*in the Proceedings of USENIX Security Symposium, 2023*

**Fortifying Federated Learning against Membership Inference Attacks via Client-level Input Perturbation**
Yuchen Yang, **Haolin Yuan**, Bo Hui, Neil Gong, Neil Fendley, Philippe Burlina, and Yinzhi Cao
*in the Proceedings of the Annual IEEE/IFIP International Conference on Dependable Systems and Network (DSN), 2023*

**ImageAlly: A Human-AI Hybrid Approach to Support Blind People in Detecting and Redacting Private Image Content**
Zhuohao Zhang, Smirity Kaushik, JooYoung Seo, **Haolin Yuan**, Sauvik Das, Leah Findlater, Danna Gurari, Abigale Stangl, Yang Wang
*in the Proceedings of the 19th Symposium on Usable Privacy and Security (SOUPS), 2023*

**Addressing Heterogeneity in Federated Learning via Distributional Transformation**
**Haolin Yuan***, Bo Hui*, Yuchen Yang*, Philippe Burlina, Neil Gong, Yinzhi Cao.
*in the Proceedings of European Conference on Computer Vision (ECCV), 2022.*

**Practical Blind Membership Inference Attack via Differential Comparisons**
**Haolin Yuan***, Bo Hui*, Yuchen Yang*, Philippe Burlina, Neil Gong, Yinzhi Cao. *: equally contributed
*in the Proceedings of Network & Distributed System Security Symposium (NDSS), 2021*

**WebAlly: Making Visual Task-based CAPTCHAs Transferable for People with Visual Impairments.**
Zhuohao Zhang, Zhilin Zhang, **Haolin Yuan**, Nata M Barbosa, Sauvik Das, Yang Wang
*in the Proceedings of Symposium on Usable Privacy and Security(SOUPS), 2021*

## Research Experience

**Model Fairness on Medical Image Processing** — Jan.2022 – now
*Johns Hopkins University* — *Baltimore, MD*

- Designed a novel image preprocessing method for deep learning models trained on medical images in both segmentation and classification tasks
- Implemented state-of-the-art semantic and medical-image segmentation works

**Practical Blind Membership Inference Attack via Differential Comparison** — Mar.2020 – Aug.2020
*Johns Hopkins University* — *Baltimore, MD*

- Implemented most of state-of-the-art membership inference attacks and defenses, such as Top-3 NN attack, Top1-threshold attack, Label only attack, etc.
- Designed a novel algorithm for the attack mechanism using differential comparison
- Designed different settings that closely simulate different environments for MI attacks

- Improved the attack performance by 20% compared to state-of-the-art MI attacks

**WebAlly–A case study of Web-task friend sourcing in solving CAPTCHA** 2020 – 2021
*University of Illinois at Urbana-Champaign* *Champaign, IL*

- Designed the privacy-guaranteed tool that utilizes friend sourcing to help people with visual impairment to solve online CAPTCHA tasks
- Empolyed DNN models to do privacy detection in given images for potential functionalities
- Implemented YOLOv3 and Microsoft Azure to compare their performances in detecting private contents

**Addressing Heterogeneity in Federated Learning via Distributional Transformation** 2021 – 2022
*Johns Hopkins University* *Baltimore, MD*

- Designed a train and test-time data transformation for heterogeneous data in FL setting
- Designed a double-input-channel model structure
- Improve FL performance (i.e., model accuracy) for clients with data of different distributional heterogeneity.

## Professional Services

- Annual Computer Security Applications Conference (ACSAC) 2023, External reviewer
- The ACM Conference on Computer and Communications Security (CCS) 2022, External reviewer
- IEEE Computer Security Foundations Symposium (CSF) 2022, External reviewer
- IEEE International Conference on Distributed Computing Systems (ICDCS) 2022, External reviewer

## Teaching/Research Assistant Experience

**Teaching Assistant** | *Johns Hopkins University* Jan.2023 – May 2023
Course: Language-based Security *Advisor: Porf. Yinzhi Cao*

**Research Assistant** | *Johns Hopkins University* Mar.2020 – Dec.2021
Department: Computer Science Department *Supervisor: Porf. Yinzhi Cao*

**Course Assistant** | *Johns Hopkins University* Sep.2020 – Dec.2020
Course: Web Security *Advisor: Porf. Yinzhi Cao*

**Teaching Assistant** | *Brandeis University* Sep.2017 – Dec.2017
Course: Precalculus Mathematics *Advisor: Porf. Rebecca Torrey*