Yiren (Aaron) Zhao

yaz21@cam.ac.uk

UK Contact: (0044) 07547842218

Education

University of Cambridge, Cambridge, UK *Grad. 2022*

PhD in Computer Science

University of Cambridge, Cambridge, UK *Grad. 2017*

Master of Philosophy in Advanced Computer Science, **Award of Distinction**

Imperial College London, London, UK *Grad. 2016*

Bachelor of Engineering in Electrical & Electronic Engineering, **First Class Honors**

Selected Awards and Honors

Research Fellowship at St John's College, University of Cambridge *2021*

This prestigious Fellowship (up to four years) offers an opportunity to carry out independent research with financial support from the College.

Apple Scholar in AI and ML *2020*

A global program created to recognize the contributions of emerging leaders in computer science and engineering at the graduate level. Received a fellowship award of around 120,000 USD from Apple Inc..

EPSRC International Doctoral Studentship joint University of Cambridge Computer Laboratory Qualcomm Premium Scholarship *2017*

Fully funded PhD scholarship for 3.5 years.

Willis Jackson Medal and Prize *2016*

For excellence in academic performance, one award per academic year.

Experience

Department of Electrical and Electronic Engineering and Imperial X, Imperial College London, Assistant Professor *2022 - Now*

- Lead a team looking at problems at the intersections between ML, hardware and security.
- Supervise Undergraduate, Master and PhD students.

Department of Computer Science, University of Cambridge, Visiting Researcher *2022 - Now*

- A visiting researcher in both the Computer Architecture and ML group.
- Supervise Undergraduate, Master and PhD students.

Department of Computer Science, University of Cambridge, Research Fellow *2021 - 2022*

- Research on how Machine Learning on unstructured, complex data types, and the potential implications to hardware systems.
- Plan and write research grant applications. Supervise Master and PhD students.

Apple AI and ML, Part-time Contractor *Dec 2021 - June 2022*
Manager: Xin Wang and Francesco Rossi

- Direct report to Apple ML compiler (middleware) team lead.
- Investigate research-orientaged projects.

Microsoft Research New England, Part-time Contractor *June 2019-Feb 2020*
Manager: Dr. Nicolo Fusi

- Working on Project AutoML
- Design novel hardware-aware AutoML methods

Microsoft Research New England, Research Intern *June 2019-Oct 2029*
Supervisor: Dr. Nicolo Fusi

- Working on Project AutoML

Microsoft Research Redmond, Research Intern *June 2018-Oct 2018*
Supervisor: Dr. Daniel Lo and Dr. Eric Chung

- Working on Project Catapult and Project Brainwave. Design and implement novel compression techniques for neural network training
- Filed several US patents for the proposed novel DNN compression techniques

Microsoft Research Cambridge, Research Intern *June 2017-Oct 2017*
Supervisor: Dr. Hitesh Ballani

- Working on Project Sirius. Design and research on new routing methods between racks in a datacenter

Papers & Thesis

DAdaQuant: Doubly-adaptive quantization for communication-efficient Federated Learning
R Honig, Y Zhao and R Mullins
International Conference on Machine Learning 2022 (ICML 2022)

Rapid Model Architecture Adaption for Meta-Learning
Y Zhao, X Gao, I Shumailov, N Fusi and R Mullins
in submission, 2021

Markpainting: Adversarial Machine Learning meets Inpainting
D Khachaturov,I Shumailov, Y Zhao, D Bates, N Papernot, R Mullins and R Anderson
International Conference on Machine Learning 2021 (ICML 2021)

Manipulating SGD with Data Ordering Attacks
I Shumailov, Z Shumaylov, D Kazhdan, Y Zhao, N Papernot, M A Erdogdu, R Anderson
Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021)

Learned Low Precision Graph Neural Networks
Y Zhao*, D Wang*, D Bates, R Mullins, P Lio and M Jamnik
The 1st Workshop on Machine Learning and Systems (EuroMLSys 2021)

Sponge Examples: Energy-Latency Attacks on Neural Networks
I Shumailov*, Y Zhao*, D Bates, N Papernot, R Mullins and R Anderson
6th IEEE European Symposium on Security and Privacy (EuroS&P 2021)

Probabilistic Dual Network Architecture Search on Graphs
Y Zhao*, D Wang*, X Gao, R Mullins, P Lio and M Jamnik
Deep Learning on Graphs: Methods and Applications (*Best student paper award*, DLG-AAAI'21)

Towards Certifiable Adversarial Sample Detection
I Shumailov*,, Y Zhao*, R Mullins and R Anderson
13th ACM Workshop on Artificial Intelligence and Security 2020 (AISEC 2020)

Pay Attention to Features, Transfer Learn Faster CNNs
K Wang*, X Gao*, Y Zhao, X Li, D Dou, X Gao, and C Xu
International Conference on Learning Representations 2020 (ICLR 2020)

Blackbox Attacks on Reinforcement Learning Agents Using Approximated Temporal Information
Y Zhao*, I Shumailov*, C Han, X Gao, R Mullins and R Anderson
IEEE International Conference on Dependable Systems and Networks Workshops (DSN-W 2020)

Focused Quantization for Sparse DNNs
Y Zhao*, X Gao, R Mullins and C Xu
Thirty-third Conference on Neural Information Processing Systems (NeurIPS 2019)

Sitatapatra: Blocking the Transfer of Adversarial Samples
I Shumailov*, X Gao*, Y Zhao*, R Mullins, R Anderson and C Xu
in submission

Automatic Generation of Multi-precision Multi-arithmetic CNN Accelerators for FPGAs
Y Zhao*, X Gao*, X Guo*, J Liu, E Wang, R Mullins, P Cheung, G Constantinides and C Xu
International Conference on Field Programmable Technology (ICFPT 2019)

The Taboo Trap: Behavioural Detection of Adversarial Samples

I Shumailov*, Y Zhao*, R Mullins and R Anderson
in submission

Characterizing Sources of Ineffectual Computations in Deep Learning Networks
M Nikolic, M Mahmoud, Y Zhao, R Mullins and A Moshovos
International Symposium on Performance Analysis of Systems and Software 2019 (ISPASS 2019)

Dynamic Channel Pruning: Feature Boosting and Suppression
X Gao*, Y Zhao*, R Mullins and C Xu
International Conference on Learning Representations 2019 (ICLR 2019)

To compress or not to compress: Understanding the Interactions between Adversarial Attacks and Neural Network Compression
Y Zhao*, I Shumailov*, R Mullins and R Anderson
The Conference on Systems and Machine Learning 2019 (SysML 2019)

Mayo: A Framework for Auto-generating Hardware Friendly Deep Neural Networks
Y Zhao*, X Gao*, R Mullins and C Xu
2nd International Workshop on Embedded and Mobile Deep Learning (Workshop of Mobisys) (EMDL 2018)

Redundancy-Reduced MobileNet Acceleration on Reconfigurable Logic For ImageNet Classification
J Su, J Faraone, J Liu, Y Zhao, D Thomas, P Leong and P Cheung
14th International Symposium on Applied Reconfigurable Computing (ARC 2018)

An Efficient Implementation of Online Arithmetic
Y Zhao, J Wickerson and G Constantinides
2016 International Conference on Field-Programmable Technology (ICFPT 2016)

Improving Compression Pipelines For Convolutional Neural Networks
Master thesis

* indicates equal contribution

Patents

Neural Network Activation Compression with Narrow Block Floating-point
D Lo, A Phanishayee, E S Chung, Y Zhao and R Zhao
US patent, US20200210838, 2020

Neural Network Activation Compression with Non-uniform Mantissas
D Lo, A Phanishayee, E S Chung, Y Zhao
US patent, US20200242474, 2020

Neural network Activation Compression with Outlier Block Floating-point
D Lo, A Phanishayee, E S Chung, Y Zhao and R Zhao
US patent, US20200210839, 2020

Adjusting Acitvation Compression for Neural Network Training
D Lo, B D Rouhani, E S Chung, Y Zhao, A Phanishayee and R Zhao
US patent, US20200264876, 2020

Teaching and Supervisions

Efficient Adversarial Training (Final year project) *2021-2022*
Maximilian Kaufmann, jointly supervised with Ilia Shumailov

Backdoors in Neural Networks (Final year project) *2021-2022*
Mikel Bober, jointly supervised with Ilia Shumailov

Reducing communication costs in Federated Learning (Final year project) *2020-2021*
Robert Honig, jointly supervised with Prof. R Mullins

Hardware-informed differentiable neural architecture search (Master project) *2019-2020*
Karl Otness, jointly supervised with Prof. R Mullins

Replay Attacks on Reinforcement Learning (Final year project) *2019-2020*
Timothy Lazarus, jointly supervised with Prof. R Mullins and Ilia Shumailov