

Yanzhao Wu

266 Ferst Drive, Room 3337, Atlanta, Georgia, 30332, USA
yanzhaowu@gatech.edu • +1 (404) 279-2853 • <http://yanzhaowu.me/>

EDUCATION

Georgia Institute of Technology, Atlanta, Georgia, USA

- Ph.D. student in Computer Science Aug 2017 – May 2021 (expected)
 - Area: **Systems & Machine Learning**
 - Focus: Big Data & Deep Learning Systems
 - Cumulative GPA: 3.91 / 4.00

University of Science and Technology of China (USTC), Hefei, Anhui, China

- Bachelor of Computer Science and Technology Sep 2013 – Jul 2017
 - Graduated with **Honors**.
 - Cumulative GPA: 3.80 / 4.30

PUBLICATION

- **Yanzhao Wu**, Ling Liu. “Selecting and Composing Learning Rate Policies for Deep Neural Networks” (Under submission)
- Ka-Ho Chow, Ling Liu, Emre Gursoy, Stacey Truex, Wenqi Wei and **Yanzhao Wu**. “Understanding Object Detection Through An Adversarial Lens.” (Accepted by ESORICS 2020)
- Wenqi Wei, Ling Liu, Margaret Loper, Ka Ho Chow, Emre Gursoy, Stacey Truex, **Yanzhao Wu**. “Cross-layer Strategic Ensemble Defense against Adversarial Examples.” In 2020 International Conference on Computing, Networking and Communications (ICNC), pp. 456-460. IEEE, 2020. (**IEEE ICNC 2020**)
- **Yanzhao Wu**, Ling Liu, Juhyun Bae, Ka-Ho Chow, Arun Iyengar, Calton Pu, Wenqi Wei, Lei Yu, Qi Zhang. “Demystifying Learning Rate Policies for High Accuracy Training of Deep Neural Networks.” In 2019 IEEE International Conference on Big Data (Big Data), pp. 1971-1980. IEEE, 2019. (**BigData’19**)
- Ka-Ho Chow, Wenqi Wei, **Yanzhao Wu**, Ling Liu. “Denoising and Verification Cross-Layer Ensemble Against Black-box Adversarial Attacks.” In 2019 IEEE International Conference on Big Data (Big Data), pp. 1282-1291. IEEE, 2019. (**BigData’19**)
- Ling Liu, Wenqi Wei, Ka-Ho Chow, Margaret Loper, Emre Gursoy, Stacey Truex, **Yanzhao Wu**. “Deep Neural Network Ensembles against Deception: Ensemble Diversity, Accuracy and Robustness” In 2019 IEEE 16th International Conference on Mobile Ad Hoc and Sensor Systems (MASS), pp. 274-282. IEEE, 2019. (**IEEE MASS 2019**)
- **Yanzhao Wu**, Ling Liu, Calton Pu, Wenqi Cao, Semih Sahin, Wenqi Wei, Qi Zhang. “A Comparative Measurement Study of Deep Learning as a Service Framework” (Accepted by **IEEE Transactions on Services Computing**.)
- Ling Liu, Wenqi Cao, Semih Sahin, Qi Zhang, Juhyun Bae, **Yanzhao Wu**. “Memory Disaggregation: Research Problems and Opportunities” In 2019 IEEE 39th International Conference on Distributed Computing Systems, pp. 1664-1673. IEEE, 2019. (**ICDCS’19**)
- Wenqi Wei, Ling Liu, Stacey Truex, Lei Yu, Mehmet Emre Gursoy, **Yanzhao Wu**. “Adversarial Examples in Deep Learning: Characterization and Divergence” (Under submission)
- **Yanzhao Wu**, Wenqi Cao, Semih Sahin, and Ling Liu. “Experimental Characterizations and Analysis of Deep Learning Frameworks” In 2018 IEEE International Conference on Big Data, pp. 372-377. IEEE, 2018. (**BigData’18**)
- Ling Liu, **Yanzhao Wu**, Wenqi Wei, Wenqi Cao, Semih Sahin, and Qi Zhang. “Benchmarking Deep Learning Frameworks: Design Considerations, Metrics and Beyond.” In 2018 IEEE 38th International Conference on Distributed Computing Systems, pp. 1258-1269. IEEE, 2018. (**ICDCS’18**)
- Wenqi Wei, **Yanzhao Wu**, Ling Liu. “DeepEyes: Integrating Deep Learning and Crowd Sourcing for Localization” (Southern Data Science Conference 2018 Research Track Poster).
- Pengcheng Wang, Jeffrey Svajlenko, **Yanzhao Wu**, Yun Xu and Chanchal K. Roy. “CCAligner: a token based large-gap clone detector” In Proceedings of the 40th International Conference on Software Engineering, pp. 1066-1077. ACM, 2018. (**ICSE’18**)

**RESEARCH
EXPERIENCE**

Pipeline Parallelism for Deep Learning Recommendation Models

- AI System SW/HW Co-Design Team, **Facebook Research** May 2020 – Aug 2020 (expected)
 - Mentor: Dheevatsa Mudigere
 - Focus: Deep Learning, **Pipeline Parallelism**
 - Goal: Apply pipeline parallelism into Facebook deep learning recommendation models to accelerate distributed recommendation model training.

High Accuracy and Robust Ensemble of Deep Neural Networks

- Distributed Data Intensive Systems Lab, Georgia Tech Aug 2019 – May 2020
 - Supervisor: Prof. Ling Liu
 - Focus: Deep Learning, **Edge AI**
 - Goal: Design and implement an ensemble framework for improving deep neural network accuracy and optimizing inference robustness on GPUs and edge devices.

Semi-automatic Hyper-parameter Tuning for Training Deep Neural Networks

- Distributed Data Intensive Systems Lab, Georgia Tech Aug 2018 – Aug 2019
 - Supervisor: Prof. Ling Liu
 - Focus: Deep Learning, Hyper-parameter Tuning
 - Goal: Design and implement a learning rate tuning system for improving accuracy and training efficiency.
 - Achievement: **LRBench**; Two papers under submission.

Experimental Analysis and Optimization of Deep Learning Frameworks

- Distributed Data Intensive Systems Lab, Georgia Tech Aug 2017 – Aug 2018
 - Supervisor: Prof. Ling Liu
 - Focus: Deep Learning Frameworks, Performance Analysis
 - Goal: Analyze the hyper-parameters and basic components of Deep Learning and optimize Deep Learning Frameworks by tuning data-related and hardware-related parameters.
 - Achievement: **GTDLBench**; Papers published in ICDCS'18, BigData'18, IEEE TSC.

A Performance Study of Deep Learning with the High-performance Storage System

- Storage Systems Research Group, **IBM Research** May 2019 – Jul 2019
 - Mentors: Dr. Daniel Waddington, Dr. Luna Xu
 - Focus: Storage Systems, Deep Learning Frameworks
 - Achievement: Conducted a comprehensive performance analysis of the high-performance storage system with different storage backends, such as **persistent memory** and SSD, with popular deep learning workloads.

Accelerating Deep Learning with Direct-to-GPU Storage

- Storage Systems Research Group, **IBM Research** May 2018 – Aug 2018
 - Mentors: Amit Warke, Dr. Daniel Waddington
 - Focus: Storage Systems, Deep Learning Frameworks
 - Achievement: Integrated the Direct-to-GPU storage system into Caffe to obtain **over 2×** performance improvement by reducing the overhead of data transmission.

DeepEyes: A Deep Learning Powered Localization System with Multi-modal Sensors

- Distributed Data Intensive Systems Lab, Georgia Tech Aug 2017 – May 2017
 - Supervisor: Prof. Ling Liu
 - Focus: Localization, Deep Learning
 - Achievement: Implemented an out-door/in-door localization system **without requiring the common localization infrastructure**, such as GPS, cellular network, and WiFi, with the help of deep learning models. A research track poster published in SDSC'18.

Detecting Large-gap Code Clones

- National High-Performance Computing Center (Hefei), USTC Sep 2015 – Jul 2017
 - Supervisor: Prof. Yun Xu
 - Focus: Source Code Processing & Indexing, Edit Distance, Detection Algorithms
 - Achievement: *CCAligner: a token based large-gap clone detector (ICSE'18)*.

Summer Research Internship on Automatic Verification

- School of Computer Science, University of Birmingham Jul 2016 – Aug 2016
 - Supervisor: Prof. David Parker
 - Focus: *LTS* (Labeled Transition Systems) Model Checker, *Game* Model Checker
 - Achievement: Implemented *LTS* model checker and *Game* model checker for PRISM, a widely applied probabilistic model checker for analysis of systems, to enable it to support non-probabilistic models further.

PEER REVIEW

- Conference: ICDE 2018, UCC 2018, BDCAT 2018, ICDCS 2019
- Journal: IEEE TKDE, ACM TOIT

OPEN-SOURCE PROJECT

- EnsembleBench: A set of tools for building good ensemble model teams for machine learning and deep learning models. (URL: <https://github.com/git-disl/EnsembleBench>)
- LRBench: A semi-automatic learning rate tuning tool to improve the deep neural network training efficiency as well as its accuracy. (URL: <https://github.com/git-disl/LRBench>)
- GTDLBench: A performance benchmark of deep learning frameworks to measure and optimize mainstream deep learning frameworks. (URL: <https://git-disl.github.io/GTDLBench/>)
- Comanche: Accelerating deep learning with Direct-to-GPU storage with a modified Caffe and DeepBench. (URL: <https://github.com/IBM/comanche>)
- CCAAligner: A token based code clone detector for detecting large-gap copy-and-paste source codes. (URL: <https://github.com/PCWcn/CCAAligner>)
- PRISM: Design and implement the LTS and Game model checker for PRISM, a widely applied model checker for system analysis. (URL: <http://www.prismmodelchecker.org/>)

SKILL

- Programming Skills: C, C++, Python, JavaScript, Java, Go, R, OpenMP, MPI, CUDA, SQL
- Machine Learning: TensorFlow, Caffe, PyTorch, Torch, MXNet, Scikit-learn, Numpy
- Computer Vision: Image Classification, Video Detection, Object Detection, OpenCV
- Big Data Analytics: Hadoop, Spark
- OS Development: Proficient with Linux and mobile OS development
- Useful Tools: Eclipse, Jupyter Notebook, Matlab, LLVM, Git, Subversion, PRISM, \LaTeX