Zhenyu Li

Homepage: https://jhsuyu.github.io/ Email: vnr3ne@virginia.edu
Mobile: +1-858-319-7681

EDUCATION

- **University of Virginia** USA
  Ph.D. in Computer Science; GPA: 4.0/4.0 *Jan 2024 – Present*
  Advisor: Prof. Chang Lou
  Research focus: Reliability of cloud-scale distributed systems

- **University of California San Diego** USA
  M.S. in Computer Science; GPA: 4.0/4.0 *Sep 2022 – Dec 2023*

- **Nanjing University** China
  B.Eng. in Software Engineering; GPA: 3.8/4.0 *Sep 2018 – Jun 2022*

PUBLICATIONS

- **Zhenyu Li**, Angting Cai, Chang Lou. "Pilot Execution: Simulating Failure Recovery In Situ for Production Distributed Systems." In Proceedings of the 23rd USENIX Symposium on Networked Systems Design and Implementation (NSDI 2026).
- Yunlong Mao, Zexi Xin, **Zhenyu Li**, et al. "Secure Split Learning against Property Inference, Data Reconstruction, and Feature Space Hijacking Attacks." In Proceedings of the 28th European Symposium on Research in Computer Security (ESORICS 2023).

INDUSTRY EXPERIENCE

- **Amazon** USA
  Software Engineer Intern *Jun 2023 - Sep 2023*
  - Built a **serverless streaming pipeline** using **Lambda** and **Kinesis** to synchronize **DynamoDB** data lake subscriptions with **ElasticSearch**, processing **10K+ events/sec** in real-time.
  - Achieved less than **100ms indexing latency**, enabling instant search capabilities across **millions of records** for downstream analytics workflows.

- **Alibaba** China
  Software Engineer Intern *Jun 2021 - Sep 2021*
  - Designed multi-threaded pub/sub system using **Kafka** with fault-tolerant message delivery for DingTalk's enterprise notification service, implementing concurrent processing to handle **1+ MB/s** throughput with exactly-once semantics, ensuring critical alerts reach all team members.

RESEARCH EXPERIENCE

- **Lift Lab, University of Virginia** USA
  Research Assistant (Advisor: Prof. Chang Lou) *Sep 2023 - Present*

- Designed **PILOT**, an execution model that simulates **recovery actions** in production environments using **phantom threads**—ephemeral, in-process execution units that rehearse recovery paths without affecting production threads.
- Implemented **context propagation** mechanism inspired by **distributed tracing** to track pilot execution across system boundaries, enabling end-to-end recovery path simulation.
- Developed **copy-on-write state shadowing** and **I/O redirection** techniques to isolate side effects, ensuring pilot execution remains isolated from production environment.
- Deployed PILOT to **5 production systems** (Solr, HDFS, Cassandra, HBase, YARN), detecting **85% (17/20)** of real-world recovery failures and discovering a **critical P1-priority bug** in the latest HBase version confirmed by the developers.

- **Nanjing University** China
  Research Intern (Advisors: Prof. Sheng Zhong, Prof. Yunlong Mao) *Mar 2021 - May 2022*
  - Designed **R^{3}eLU**, a privacy-preserving activation function for **split neural networks** combining **randomized response** and **Laplace mechanism**, with **dynamic privacy budget allocation** based on neuron importance for vertically partitioned collaborative learning.
  - Achieved (ϵ, δ)**-differential privacy** for both host and guest parties while maintaining **93-95% model accuracy** on MNIST and reducing **property inference attack success** from **80%+ to 50-60%**.

PROGRAMMING SKILLS

- **Languages:** Java, Python, C/C++, Go, CUDA
- **Systems:** HBase, HDFS, YARN, Solr, Cassandra, Kafka, etcd

AWARDS & HONORS

- UVA Outstanding Teaching Assistant Award 2025
- SOSP 2024 Travel Grant 2024