

JINGYU "NORA" TANG

GHC '23 | LinkedIn | Personal Portfolio | +1(608)-733-7355 | jingyu_tang@alumni.brown.edu | Austin, TX

EDUCATION

Brown University

Sc.M. in Computer Science; GPA: 4.0/4.0

T.A. Experience: CSCI 1260 Compilers and Program Analysis, CSCI 1952Q Algorithmic Aspects of Machine Learning

Relevant Courses: Advanced Database, Privacy Systems, Network, Compiler, Computer Systems, Deep Learning

Sep. 2022 – May 2024

Providence, RI

Case Western Reserve University

B.A. in Computer Science and Mathematics; GPA: 3.94/4.0

Awards: The Webster Godman Simon Mathematics Award; Summa Cum Laude

Relevant Courses: Software Engineering, Database, Computer Security, Data Mining, Machine Learning

Aug. 2019 – May 2022

Cleveland, OH

WORK EXPERIENCE

SDE Intern spec. in Data Engineering

Expedia Group

May 2023 – July 2023

Austin, TX

- Constructed and integrated real-time and batch concurrent data pipelines, enabling daily and weekly analysis of approximately **75k real-time booking events** as the source-of-truth for EG-wide business intelligence operations.
- Collaborated actively with the core DevOps team of 5, adhering to **Agile Scrum** methodology, to conclude the design and develop the architecture, schema, UML, etc.
- Implemented the processor with Spring Boot, achieving **94% unit test coverage**, and persisted data in Elasticsearch and Amazon S3 buckets.
- Deployed Apache Kafka topology, schema, streams, producers and consumers to consume incoming data and publish processed events with **latency below 50ms**.

SDE Intern

Industrial and Commercial Bank of China

July 2022 – Aug. 2022

Remote

- Designed a full-stack data lineage application, supporting data governance with approximately **20k daily events**, focusing on parsing SQL tasks into a visually intuitive lineage map at the metadata level.
- Engineered a robust tech stack, including Spring Boot as the framework, Druid as the SQL tasks monitor, Druid SQL Parser as the SQL and lineage parsing tool, and Neo4j as the **lineage database management**.
- Programmed an interactive front-end with **data visualization** capabilities, employing Vue.js and Neovis for an engaging user experience.

PROJECTS

DUSE: A Distributed Search Engine for USENIX Paper | JavaScript

Jan. 2024 – May 2024

- Collaborated with a team of 4 to implement a distributed search engine for the USENIX publication repository, capable of handling searches **over 100k words**.
- Utilized the map-reduce model, consistent hashing, and Remote Procedure Call communication to achieve scalability and fault-tolerance.
- Deployed the search engine on **Elastic Compute Cloud (EC2)** nodes.

Adversarial Learning Is All Semantic Segmentation Needs | Python

Jan. 2024 – May 2024

- Applied adversarial learning combined with cross-entropy loss to improve semantic segmentation performance, using a dataset of **5108 aerial images** to identify forested and non-forested areas.
- Devised a model with a segmentation network and a fully convolutional discriminator network over 800 epochs.
- Achieved an average **loss of 0.7** and an average **mean IoU of 80%**, demonstrating the effectiveness of the model with a relatively small dataset.

RFC-compliant IP/TCP Protocols | Go

Sep. 2023 – Dec. 2023

- Assembled a virtual link layer interface with UDP sockets, network, and transport layers, achieving the transmission of a **1 megabyte** data over a lossy connection with a **2% drop rate**.
- Enforced key RFC-compliant TCP functionalities such as the sliding window protocol, and retransmission.

Distributed Blockchain Protocol | Go

Jan. 2023 – May 2023

- Developed a fast and scalable blockchain server embedded with a distributed system, implementing a distributed system to facilitate seamless trading across multiple users.
- Established a robust and thread-safe server and database using Go, ensuring the safety and promptness of trading operations.

SKILLS

Programming Languages: Java, Go, JavaScript, Python, OCaml, SQL, C

Frameworks and Tools: Spring Boot, Git, Apache Spark, Apache Kafka, Elasticsearch, Amazon S3, Datadog, GitHub Actions, Splunk